Individual and Institutional Investor Behavior

by

Noah S. Stoffman

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in the University of Michigan
2008

Doctoral Committee:

Associate Professor Tyler G. Shumway, Chair
Professor Miles S. Kimball
Associate Professor Uday Rajan
Assistant Professor Feng Li
Assistant Professor Kathy Z. Yuan

For my parents

Acknowledgements

The work presented here would not have been possible without the help of my teachers, colleagues, family, and friends.

In particular, many helpful discussions with Sreedhar Bharath, Bob Dittmar, Ryan Israelsen, Amiyatosh Purnanandam, Amit Seru, Sophie Shive, and Clemens Sialm are gratefully acknowledged.

I thank each of the members of my dissertation committee for their time and many helpful suggestions. I have particularly benefited from my numerous interactions with Uday Rajan and Kathy Yuan.

I am especially indebted to Tyler Shumway, who not only guided me through the dissertation process, but has also become a good friend. The first chapter of this dissertation was written jointly with him and Amit Seru, and I am grateful to both for their collaboration.

I also thank Jussi Keppo, who helped obtain the data set used throughout this dissertation, and explained some of the idiosyncracies of the Finish stock market.

Finally, I am eternally grateful to Jennifer Richler, without whose love and support I would never have been able to complete this dissertation.

Table of Contents

List of Tables

List of Figures

Chapter I

Introduction

Academic research has recently documented a wide range of behaviors among individuals that is difficult to reconcile with the classical view of a rational *homo economicus*. In contrast to the "smart money" controlled by professional money managers, it seems plausible that individual investors may consistently suffer from costly behavioral biases. These biases may affect not only the performance of individual investors, but also aggregate to cause price distortions. Given the critical role for prices in capital allocation, such distortions could be economically important.

The three papers presented in this dissertation explore different aspects of this issue. Chapter I examines whether individual investors learn from their trading experience to avoid a particular behavior bias and improve their returns. If individual investors suffer from behavioral biases but quickly learn to avoid them as they become more experienced, then the average degree of bias in the population will diminish without a continuous influx of new, bias-prone investors. Previous estimates of the speed of learning suggest that biases would quickly die out in the population, but the results presented here suggest that in fact the speed at which individuals learn is much slower. The relatively slow speed at which investors learn suggests that we cannot rule out the possibility that individual investors suffering from behavioral biases could distort asset prices.

Given these results, a natural question to ask is whether individual investors actually affect stock prices. This is explored in Chapter II. In particular, I examine in which direction prices move when trading occurs between two individuals, between individuals and institutions, or between two institutions. I find that prices move in the direction of institutional trading. That is, when individuals buy shares from institutions, prices fall; and when individuals sell shares to individuals, prices rise. These results are consistent with the notion that individual investors supply liquidity to institutions, and not that individuals actively move prices.

In Chapter III, I implement a strategy for identifying those trades placed by individual investors that are particularly likely to be based on information. I then examine whether these trades earn higher *ex post* returns, and I find that they do. This indicates that while many individual investors may generate poor investment performance, at least some individual investors are able to trade on information. The technique presented here could also be used to estimate the level of informed trading that occurs among individuals at any particular time.

1.1 Data

In each of the studies presented here, I use a data set that includes the complete trading records of all individual and institutional investors in Finland over a nine-year period. The remarkable richness of these data allow a uniquely detailed examination of the behavior of individual and institutional investors in a financial market. The data come from the central register of shareholdings in Finnish stocks maintained by Nordic Central Securities Depository (NCSD), which is responsible for the clearing and settlement of trades in Finland. Finland has a direct holding system, in which individual investors' shares are held directly with the CSD. Since our data come from the CSD, they reflect the official record of holdings and are therefore of extremely high quality. In particular, shares owned by individuals but held in street name by a brokerage firm are identified as belonging to the individual, and shares for each individual are aggregated across brokerage accounts, regardless of whether they are held in street name. This allows a clean identification of which type of investor owns shares. Grinblatt and Keloharju (2000, 2001a, 2001b) use a subset of the same data, comprising the first two years of our sample period.[1] The data include the transactions of nearly 1.3 million individuals and firms, beginning in January, 1995 and ending in December, 2003. Summary statistics relevant to the analysis in each chapter are presented separately.

While our dataset includes exchange-traded options and certain irregular equity securities, we focus on trading in ordinary shares. Trading in Finland is con-

[1] These references provide a detailed discussion of the data.

ducted on the Helsinki Stock Exchange, which is owned by OMX, an operator of stock exchanges in Nordic and Baltic countries. Trading on the Helsinki exchange begins with an opening call from 9:45–10:00 a.m., and ends with a closing call from 6:20–6:30 p.m. Continuous trading during regular hours is conducted through a limit order book.

The transaction data include the number of shares bought or sold, corresponding transaction prices, and the trade and settlement dates, although trades are not time-stamped. Additional demographic data, such as the account-holder's age, zip code, and language are also included, as are initial account holdings at the beginning of our sample period.

Only the direct holdings and transactions of individuals are available. This means that for an individual who directly trades shares of Nokia and holds a Finnish mutual fund that owns shares of Nokia, we will observe only trades in the former. The trades of the mutual fund are included in the dataset, but are identified as holdings of the mutual fund company, and cannot be tied to the individual. However, our wealth calculations allow us to compare the importance of the individual investors as a group to that of other market participants. On average, individuals hold 12.6 percent of all equity held by Finnish investors, including financial institutions, government funds, nonprofit organizations and nonfinancial corporations. This is more than financial institutions, which hold an average of 9.6 percent during our sample period. The majority of equity is held by the government (34.7%) and nonfinancial firms (33.4%), although these investors trade relatively less and may do so for strategic reasons that are not directly linked to profit-maximization.

Chapter II

Learning By Trading

Academics have recently shown interest in the investment behavior and performance of individuals, a field that has been called 'household finance' by Campbell (2006). Over the past decade, several researchers have documented a number of behavioral biases among individual investors. More recently, researchers have found evidence that some individual investors are more informed or skilled than others.[1] Considering these findings, it is natural to ask how skilled or informed investors acquire their advantage. For example, do investors learn by trading? If so, to what extent do investors improve their ability and to what extent do they learn about their inherent ability? And how quickly do investors learn? In this chapter, we exploit trading records to study both average investor performance and the strength of the behavioral bias known as the disposition effect.[2] We corre-

[1]Coval, Hirshleifer, and Shumway (2005) document significant performance persistence among individuals. Ivković and Weisbenner (2005) find that individuals place more informed trades in stocks of companies located close to their homes, and Ivković, Sialm, and Weisbenner (2006) show that individuals with more concentrated portfolios tend to outperform those who are more diversified. Linnainmaa (2007) finds that individuals who trade with limit orders suffer particularly poor performance.

[2]The disposition effect is the propensity of investors to sell assets on which they have experienced gains and to hold assets on which they have experienced losses. The effect was

late performance and disposition with investor experience and investor survival rates to determine whether and how investors learn by trading.

Motivated by the existing economics literature on learning, we consider two specific ways in which investors can learn. First, in the spirit of classical learning-by-doing models (Arrow 1962, Grossman, Kihlstrom, and Mirman 1977), investors might improve their ability as they trade. Second, as investors trade they might realize that their inherent level of ability is low and decide to stop actively trading. This decision to either continue or stop trading is a feature of the recent learning model of Mahani and Bernhardt (2007). In our analysis we take either of these two scenarios as evidence of learning by trading. Although these types of learning are different, they are not mutually exclusive.

To clarify the model of learning we have in mind, consider the case of an individual who decides to begin trading. The investor must decide which of the myriad sources of market information and investment advice available to her to take seriously. She could consult standard news sources, internet sites, investment newsletters, and neighbors or friends. She might also consider the advice of brokers, news analysts, authors of books and magazines, and finance professors. To the extent that these sources fail to agree completely, individuals must determine how much decision weight to assign to each source. Moreover, the quality

first proposed by Shefrin and Statman (1985), and was subsequently documented in a sample of trading records from a U.S. discount brokerage firm by Odean (1998). The effect has been found in other contexts, including in Finland (Grinblatt and Keloharju 2001a), China (Feng and Seasholes 2005, Shumway and Wu 2006), and Israel (Shapira and Venezia 2001); among professional market makers (Coval and Shumway 2005), mutual fund managers (Frazzini 2006), and home sellers (Genesove and Mayer 2001); and in experimental settings (Weber and Camerer 1998). We focus on the disposition effect because it is a robust empirical finding and it is relatively easy to measure.

of these sources is likely to differ across individuals: some investors may know executives at a firm, while others will not. As investors begin trading, they can learn to which of the various sources they should pay more attention. This can be thought of as improving ability through learning. Investors who only have access to poor sources of information cannot improve by focusing more on particular sources. Instead, they will learn that they have no useful information and will stop trading actively, choosing instead to invest in a passive investment such as an index fund.[3] This is also a type of learning, but rather than improving her ability, the investor learns *about* her ability.

This example helps to elucidate the two types of learning we examine in this chapter. First, investors may learn in such a way that their ability improves, perhaps by learning which of their sources of information are particularly useful. We measure this type of learning by correlating investors' performance and susceptibility to the disposition effect with their experience. Second, investors may learn about their ability, perhaps by learning whether they have access to useful information sources. Because investors who learn that their ability is poor will cease to trade, we identify this type of learning by examining attrition in our data. Once we control for unobserved individual heterogeneity and attrition, any learning that remains will be of the first type, a direct improvement in ability. Since we track individuals over time—both their survival rates as well as their performance

[3] This example is very similar in spirit to well-known 'bandit' problems. For example, Bolton and Harris (1999) study the strategic interaction of agents in an experimentation game. In our setting, we can think of investors as being randomly assigned slot machines with different unobservable expected payoffs, and experimenting to learn what the payoffs are. The random assignment is analogous to the assignment of inherent ability.

and their level of the disposition effect—we are able to differentiate between these two types of learning. Ours is the first paper to identify and measure both types of learning.

Differentiating between types of learning allows us to determine how quickly investors learn in a meaningful way. Estimating learning without controlling for heterogeneity and attrition results in inflated improvement estimates, which do not correspond to the experience of any particular type of investor. Measuring the speed of learning in a meaningful way is important for a number of reasons. If investors learn quickly and there is low turnover in the population of investors, behavioral biases are unlikely to significantly affect asset prices. Moreover, if they learn relatively quickly then the 'excessive' trading documented by Odean (1999) and Barber and Odean (2001) may be justified, because investors may optimally choose to trade more actively if they know they will improve with experience. Finally, policy makers should consider the speed at which individuals learn when determining the costs and benefits of different trading mechanisms.

The type of learning we observe, besides being important from a theory perspective, also has significant implications for policy. For instance, if investors can significantly improve their ability from experience then investor education initiatives may be worthwhile and perhaps individual investors should be encouraged to trade. However, if the ability to invest successfully is relatively fixed, then screening mechanisms, or tests that measure and reveal inherent ability, have more value than education. Finally, thinking about both the speed and the type of learning, has implications for market efficiency. For example, if many inexperienced investors begin trading around the same time, and they learn slowly, their trading could lead to time-varying market efficiency.

We test our hypotheses with a remarkable dataset that includes the complete trading records of investors in Finland from 1995 to 2003, including more than 22 million observations of trades placed by households. We use these data to estimate disposition and calculate performance at the account level. Our disposition estimates indicate that a median individual in our sample is 2.8 times more likely to sell a stock when its price has risen since purchase than when its price has fallen. We exploit the panel structure of our data to examine whether individual investors learn to avoid the disposition effect and improve their performance as they trade. In particular, we estimate the mean return and the disposition effect for each account and year in our sample and relate these estimates to experience, past returns, and various demographic controls.

We measure investing experience with both the number of years that an investor has been trading and with the cumulative number of trades that an investor has placed. Of course, investors may gain experience by actively trading securities and observing the results of each trade. If this is the primary way in which investors learn, then cumulative trades will predict future investment performance and the disposition effect. However, investors may also learn by observing market quantities and considering the outcomes of hypothetical trades based on, for example, a particular information source. If this is the primary way that investors learn, then years of experience will be a better predictor of investment performance and the disposition effect than cumulative trades.[4]

Our tests provide robust evidence of learning by trading at the individual level. We show that in a simple model of learning, performance improves and the disposition effect declines as investors become more experienced. An extra year of experience is associated with an improvement in average returns of approximately 40 basis points (bp) over a 30-day horizon, and a reduction in the disposition effect of about 5 percent. However, we argue that individual heterogeneity and survivorship effects are likely to significantly affect these simple estimates, making them difficult to interpret.

We adjust for survivorship and heterogeneity with a modified Heckman selection model that allows for individual fixed effects. The Heckman selection model is a two-stage instrumental variables model that adjusts for the possibility that the composition of the sample is endogenous. As instruments, we construct two variables that satisfy the necessary exogeneity conditions—that is, they are likely to affect the probability of an investor remaining in the sample, but are unlikely to affect changes in the investor's performance or disposition effect except through their effect on survival. The first variable is an indicator for whether the individual inherited shares in the previous year due to the death of a relative. We conjecture that an individual who inherits shares is more likely to trade in the future, perhaps because their wealth has increased, or perhaps because the new

[4]It is also possible, of course, that investors learn by considering the returns to hypothetical trades before they ever start trading. If this is the only way in which investors learn then we will find no evidence of learning by trading.

shares cause them to pay more attention to the stock market. This satisfies the exogeneity condition since inheritance of shares from a relative is unlikely to directly affect changes in the performance or the disposition effect of an individual. The second variable is the variation in the returns across all positions taken by an account in the previous year. This variation is a measure of the consistency of an investor's performance, and we conjecture that an investor with more variable performance is more likely to stop trading. While an investor's previous average performance is likely to be related to her future performance, there is no reason to believe that the consistency of an investor's past performance should directly affect *changes* in her future performance or disposition effect.

In our first-stage estimates we confirm that investors with poor performance are those who are more likely to cease trading. Moreover, more successful investors continue to trade actively. We also find that adjusting for survivorship and heterogeneity in our learning reduces our learning coefficients by about one-half to three-quarters. This suggests that while individuals do learn to improve their trading ability with experience, they primarily learn about their own inherent trading ability and they cease to trade if their ability is low.

Some evidence about learning by trading exists. Feng and Seasholes (2005) give evidence that investors, in aggregate, display significantly less disposition over time, estimating that for sophisticated investors the disposition effect is essentially attenuated after about 16 trades. These estimates do not appear to be consistent with the findings of Frazzini (2006), which shows that mutual fund managers, who trade substantially more than most individuals, display a significant disposition effect. Furthermore, since Feng and Seasholes (2005) do not adjust for heterogeneity and attrition, it is impossible to tell whether their estimates imply that a particular investor who trades 16 times will no longer exhibit the disposition effect or whether they imply that most investors who exhibit the disposition effect will cease trading before they place 16 trades. Nicolosi, Peng, and Zhu (2004) show that the trading performance of individuals appears to improve with trading experience, estimating that individuals can improve their risk-adjusted portfolio return by about two percent per year (or about 0.8 bp per day) over a 3-year period. This seems quite large. Again, it is not possible to tell whether this estimate is driven by the most successful investors surviving or by the least successful investors improving their ability.[5]

While our tests have some features in common with existing papers, they differ from the literature in a number of important respects. First, unlike other papers, our tests use measures of performance and estimates of the disposition effect that are specific to *individuals*, allowing us to track particular individuals over time. This allows us to control for investor heterogeneity and survivorship effects and allows us to separate the two types of learning as well as to estimate the speed of learning. It also ensures that each observation in almost all of our tests is an average or regression coefficient for one individual in one particular year. Thus, we can be sure that our results hold for the average active investor. Put another way, investors who trade disproportionately often do not receive disproportionate

[5] Other papers that find some learning in various settings include List (2003), Barber, Odean, and Strahilevitz (2004), Linnainmaa (2006), and Choi, Laibson, Madrian, and Metrick (2007).

weight in our estimates. Finally, using only one observation per person per year reduces the likelihood that our standard error estimates are incorrect because of correlation among our regression residuals. Given the unique features of our data and our test methods, the results of our hypothesis tests add significantly to the literature on financial learning.

The rest of the chapter is organized as follows. Section 2.1 describes the hypotheses we test and our statistical methods, while Section 2.2 provides detail on our data. Section 2.3 discusses our results and Section 2.4 concludes.

2.1 Hypotheses and Methods

We test three hypotheses in this chapter. Our first hypothesis (H1) is that individual investors learn by trading. More importantly, we are interested in understanding how investors learn by trading. Our second and third hypotheses clarify the type of learning in H1. Our second hypothesis (H2) is that investors learn about their inherent ability by trading, and our third hypothesis (H3) is that investors learn to improve their ability over time. Of course, if H1 is false, then H2 and H3 will also be false. We test H2 by examining the importance of individual heterogeneity and attrition in our sample. The notion in H2 is that individuals stop trading if they realize their ability is low. Once we control for unobserved individual heterogeneity and attrition, any learning that remains will support H3—that is, it suggests that individuals learn to directly improve their ability. To examine our hypotheses we test a number of related predictions. Our prior belief for each of these predictions is that investors both learn about their inherent ability and learn to improve their ability. This section motivates and describes our tests in more detail. It also describes some of the statistical methods of our tests.

2.1.1 Measuring performance

Investor performance is the primary variable that we correlate with experience to test our hypotheses. Measuring the performance of individual investors is a significant challenge. Our data, like others that are available, do not include all non-equity securities that may be held by an investor, so it is impossible to measure the return for the investor's entire portfolio. This is made more difficult by the fact that the amount of money an individual has invested in equities often fluctuates significantly over time. Since we cannot accurately measure portfolio returns, we measure performance by examining the average return of stocks purchased. However, this generates a new problem—comparing the returns on holding periods of different lengths. For example, it is particularly difficult to compare the performance of one investor who holds a stock for one week and earns a holding period return of 3 percent to that of another investor who holds a stock for one year and earns a holding period return of 15 percent.

Given the challenge of calculating performance, we take a straightforward approach that is nevertheless likely to capture much of the relevant information in the individual's returns. In particular, we calculate the returns earned by the purchased stock in the 30 trading days following each of an investor's purchases. We choose to examine 30-day returns because the median holding period in our data is 39 trading days, but all of our findings remain unchanged if we use a 10-, 45-, or 60-day holding period. Importantly, we truncate this calculation window at the length of the actual holding if it is shorter than 30 days. That is, the 30-day return for investor i holding stock j is,

$$R_{i,j}(t) = \frac{P_j(t + \min(s, 30))}{P_j(t)} - 1,$$

where $P_j(\cdot)$ denotes the stock's closing price adjusted for splits and dividends, t denotes the purchase date, and s denotes the actual holding period.

Our approach is an attempt to deal with the problem of comparing returns over similar holding periods while ensuring that the actual selling decisions of investors affect their performance. By measuring returns this way, we hope to capture the value of short-term signals that the investor may have received. Looking over longer horizons would introduce considerable noise into our return estimates.

To adjust for risk, some of our results use risk-adjusted returns (or alphas). We use a four-factor model to adjust returns for known risk factors. In addition to a value-weighted market return, we construct three factor-mimicking portfolios (SMB, HML, and UMD; see Fama and French (1993) and Carhart (1997)). To construct the HML and SMB factors, we augment our data with information from Thomson Financial's DataStream database. We take quarterly data on shares outstanding and book value of equity using definitions as similar as possible to those of Fama and French.

Each quarter, we sort firms into deciles independently along three dimensions: market capitalization (price per share times shares outstanding), market-to-book (price per share divided by book value of equity per share), and past returns (over months $t-12$ to $t-2$). The small-minus-big (SMB) portfolio consists of a long position in the top half of stocks by market capitalization, and a short position in the bottom half. The high-minus-low (HML) portfolio consists of a long position in those stocks in the top 30 percent of market-to-book, and short position in stocks in the lowest 30 percent. The up-minus-down (UMD) portfolio consists of a long position in stocks in the top decile of past return performance, and a short position in stocks in the bottom decile. We use the overnight interbank lending rate in Finland as a proxy for the risk-free rate. Beginning in 1999, this is equivalent to the Euribor rate.

To measure a firm's risk in a particular year, we regress daily returns in excess of the risk-free rate on the daily returns of the four factor-mimicking portfolios and a constant. To ensure that our estimates are not contaminated by well-known problems associated with nonsynchronous trading, we follow the two-stage least squares approach of Scholes and Williams (1977, pp. 316–319).

We use the estimated factor betas from these regressions to calculate risk-adjusted daily returns for each stock, and we use these adjusted returns in the same performance regressions we previously estimated with raw returns. The risk adjusted return is defined as

$$\alpha_{i,t} = R_{i,t} - \left(\beta_i^M \mathrm{RMRF}_t + \beta_i^S \mathrm{SMB}_t + \beta_i^H \mathrm{HML}_t + \beta_i^U \mathrm{UMD}_t\right), \qquad (2.1)$$

where $R_{i,t}$ denotes the raw return of stock i on date t, β_i^M is stock i's market beta, β_i^S is the beta on the size factor (SMB), β_i^H is the beta on the value/growth factor (HML), and β_i^U is the beta on the momentum factor (UMD). The returns RMRF, SMB, HML, and UMD are the returns on the market portfolio, and the three factor-mimicking portfolios, respectively. Given the daily time-series of factor-mimicking portfolios constructed as described above, we have a daily series of risk-adjusted returns, which we can then use instead of raw returns in several of our tests. As with the raw returns, we calculate holding period returns over a fixed interval of 30 days, but truncate the holding period at the length of the actual holding period.

Given our measure of performance, we perform a few tests to be sure that it has similar properties to the measures used previously, such as poor investor performance on average and persistence in performance. However, our main prediction about performance relates directly to the central hypothesis of our paper:

Prediction 1. *Individuals with more investment experience have better investment performance.*

We test this conjecture at this point by regressing our performance measure on our experience measures. Performance is measured as an investor's average return over the 30 trading days following each purchase. Our primary experience variables include the number of years that an investor has been in our data and the cumulative number of trades the investor has placed. We include a quadratic term for each experience variable to allow investors to learn more slowly over time. Notably, evidence for this prediction is consistent with investors learning by trading as predicted by H1, but it does not allow us to differentiate between H2 and H3. Later in the chapter we adjust for individual heterogeneity and potential survivorship bias. Results of the tests related to performance are presented in Section 2.3.1.

2.1.2 Measuring disposition

Another outcome variable that we track to evaluate whether investors learn with experience is the disposition effect. Previous researchers have measured the disposition effect in a number of ways. Odean (1998) compares the proportion of losses realized to the proportion of gains realized by a large sample of investors at a discount brokerage firm. Grinblatt and Keloharju (2001a) model the decision to sell or hold each stock in an investor's portfolio by estimating a logit model that includes one observation for each position on each day that an account sells any security. Days in which an account does not trade are dropped from their analysis.

As Feng and Seasholes (2005) point out, a potential problem with these and similar approaches is that they may give incorrect inferences in cases in which capital gains or losses vary over time. Hazard models, which have been extensively applied in a number of fields including labor economics and epidemiology, are ideally suited to our setting. Feng and Seasholes (2005) implement a parametric hazard model. However, they pool their data and estimate the hazard regression only once for all investors together, ignoring most individual-level heterogeneity and any survivorship bias. Since our focus is on estimating disposition at an individual level, we estimate the hazard regression for each investor and year. Implementation of the hazard model uses *all* data about the investor's trading and the stock price path, rather than just data on days when a purchase or sale is made. That is, it implicitly considers the hold-or-sell decision each day. This improves our disposition estimates, and gives us more power with which to investigate learning. More important, individual level estimates enable us to test H2 and H3 by explicitly accounting for investor attrition and heterogeneity.

We use a Cox proportional hazard model with time-varying covariates to model the probability that an investor will sell shares that he currently holds. We count every purchase of a stock as the beginning of a new position, and a position ends on the date the investor first sells part or all of his holdings. Alternative definitions of a holding period, such as first purchase to last sale, or requiring a complete liquidation of a position, do not substantively alter our results. Our time-varying covariates include daily observations of some market-wide variables and daily observations of whether each position corresponds to a capital gain or loss.

Proportional hazard models make the assumption that the hazard rate is

$$\lambda(t) = \phi(t) \exp(x(t)'\beta). \tag{2.2}$$

Here, the hazard rate, $\lambda(t)$, is the probability of selling at time t conditional on holding a stock until time t, and $\phi(t)$ is referred to as the baseline hazard. The term $\exp(x(t)'\beta)$ allows the expected holding time to depend on covariates that vary over time. In our specification, each of the covariates changes daily. Since we estimate the hazard model for each investor-year, the baseline hazard rate describes the typical holding period of just one investor in one particular year. The Cox proportional hazard model does not impose any structure on the baseline hazard, and Cox's (1972) partial likelihood approach allows us to estimate the β coefficients without estimating $\phi(t)$. The method also allows for censored observations, which is important in our setting because investors have not always closed a position by the end of a given year. Details about estimating the proportional hazard model can be found in Cox and Oakes (1984).

The particular specification we use for investor i's hazard to sell position j is

$$\lambda_{i,j}(t|x_{i,j}(t)) = \phi_i(t)\exp\left\{\beta_i^d I(R_{i,j}(t) > 0) + \beta_i^r \bar{R}_{M,t} + \beta_i^s \sigma_{M,t} + \beta_i^V V_{M,t}\right\}. \quad (2.3)$$

The key variable in this regression is $I(R_{i,j}(t) > 0)$, an indicator for whether the total return on position j from the time of purchase up until time t is positive. Investors who suffer from the disposition effect are more likely to sell when this condition is true, so they will have positive values of β_i^d. In the rest of the chapter, we refer to β_i^d as the 'disposition coefficient.' The total return variable includes any dividends or other distributions, and is calculated using closing prices on all days, including the date of purchase. We use the closing price on the purchase date instead of the actual purchase price to ensure that our results are not contaminated by microstructure effects.[6] In addition, we include as controls three 5-day moving averages of market-level variables to ensure that we are not capturing selling related to market-wide movements: market returns ($\bar{R}_{M,t}$), squared market returns ($\sigma_{M,t}$), and market volume ($V_{M,t}$). We repeat this estimation via maximum likelihood each year from 1995–2003 for each individual i who places at least seven round-trip trades in a year.

To investigate learning, we use our disposition estimates as dependent variables in cross-sectional regressions. Because these estimates are measured with noise, we use weighted least-squares (WLS) regressions where the weight given an observation is proportional to the reciprocal of the estimated variance of the disposition coefficients. The standard errors used in this calculation are calculated using the robust 'sandwich' estimator, with clustering by security.

We perform a few tests on our disposition coefficients to be sure that our measure has similar properties to the measures used previously. In particular, for investors to have an incentive to learn to avoid the disposition effect, it must be a behavioral bias that is costly to them. We examine whether disposition is a somewhat stable, predictable attribute of a particular investor and whether investors with more disposition have inferior investment performance. However, our main prediction about disposition is closely related to our central hypothesis:

[6]In addition, our data do not include transaction prices during the first three months of 1995, but we do have closing prices during this period.

Prediction 2. *Individuals with more investment experience have a weaker disposition effect.*

We test this prediction initially by regressing disposition coefficients on our experience measures. As in our performance results, we use the number of years that an investor has been in our data, and/or the cumulative number of trades to measure experience. Also similar to our earlier tests on performance, support for this prediction is consistent with H1, but it does not allow us to differentiate between H2 and H3. Again, later in the chapter we adjust for individual heterogeneity and survival to differentiate between the two types of learning.

Our empirical strategy differs from that of Feng and Seasholes (2005), who estimate one large hazard model with all investors pooled together. Their hazard model includes an experience variable, and an interaction term of $I(R_{i,j}(t) > 0)$ with their experience variable. We avoid this approach because it gives inordinate statistical weight to investors who trade frequently. To see the sort of problem that this approach might generate, suppose that investors trade according to different styles. Some investors trade like day traders, exploiting short-term supply and demand imbalances, while others trade like fundamental analysts, looking for larger and longer-term mispricing. If we aggregate traders by assuming that each of an individual's trades is a separate, conditionally independent observation, the inferences we make about learning will be driven by the types of investors who trade the most. Accordingly, we could mistakenly attribute the improvements of a select few individuals to the whole population. While we avoid this sort of problem by making an individual-year the unit of observation, we estimate a few hazard models with the Feng and Seasholes (2005) method to allow comparison with their results.

2.1.3 Survivorship and heterogeneity

Our predictions so far have ignored heterogeneity and survivorship. However, we need to consider heterogeneity and attrition to have persuasive inferences about learning and to test H2 and H3. Through wealth effects, significant heterogeneity in ability can plausibly make it appear as if there is learning when in fact there is none. More important, simple models that neglect survivorship may provide evidence consistent with investors learning, but they will not be able to determine whether poor investors actually improve their ability or simply drop out after learning about their inherent inability. To avoid these problems, we carefully control for investor heterogeneity and survivorship bias. We use two methods to disentangle the two effects. First, we control for time-invariant unobserved individual characteristics by including individual fixed effects in our learning regressions. Any additional improvement over time is then likely attributable to improvement of ability. Second, we directly examine how much ceasing to trade (or learning about ability) affects our inferences about learning by using a modified version of the selection model introduced by Heckman (1976).

The classic Heckman model involves a two-stage procedure. In the first stage, a selection model is constructed to predict which observations will be observable in the second stage. In the second stage, the regression of interest is estimated with an adjustment for survivorship bias. We modify this procedure to account for both survivorship bias and individual heterogeneity, adopting the empirical strategy of Wooldridge (1995), which modifies the Heckman model to allow for fixed effects. Intuitively, this approach accounts for survivorship by estimating the selection model every year and including the inverse Mills ratios (the conditional probability that an individual would not cease to trade) of each selection equation in the learning regression model. Individual time-invariant heterogeneity is accounted for in this method by running the learning regression in first differences. More concretely, the learning regression model we estimate is

$$\Delta y_{i,t+1} = \beta \Delta x_{i,t} + \rho_2 d2_t \lambda_{i,t} + \ldots + \rho_T dT_t \lambda_{i,t} + \epsilon_{i,t}, \tag{2.4}$$

where $\lambda_{i,t}$ are the Inverse Mills ratios from a year t cross-sectional probit model (the selection model) and $d2_t, \ldots, dT_t$ are year dummies. Including these variables in the learning regression accounts for the impact of the selection equation. Note that a joint test of $\rho_t = 0$ for $t = 2, \ldots, T$ is a test of whether survivorship bias is a concern.

The first-stage uses cross-sectional probit regressions to predict whether or not the individual ceases to trade in a given period. The probit regressions include a constant, linear and quadratic experience terms, the number of different stocks the investor trades, the individual's average return in the previous year, and the individual's average daily marked-to-market total portfolio value. As instruments, we use the following variables: (1) a dummy variable for whether an investor inherited shares in the previous calendar year; and (2) the cross-sectional standard deviation of the individual's previous-year 30-day return. As we will explain below, both these variables are likely to satisfy the necessary exogeneity conditions—that is, they are likely to affect the probability of remaining in the sample, but are unlikely to affect changes in an individual's performance or disposition effect except through their effect on survival.

For our first instrument, we conjecture that an individual who inherits shares is more likely to trade in the future, perhaps because their wealth has increased, or because the new shares cause them to pay more attention to the stock market. This satisfies the exogeneity condition since inheritance of shares from a relative is unlikely to directly affect changes in the performance or disposition effect of an individual. In the data, when an investor dies and shares are transferred to an heir, it appears as a transaction with the account of the deceased selling shares and the account of the heir purchasing shares. A special code identifies the transaction as an inheritance. In our sample death transfers are evenly spread over the sample (ranging from 62 in 1995 to as high as 443 in 2003). Our second instrument is the variation in the returns across all positions taken by an account in the previous year. This variation is a measure of the consistency of an investor's performance, and we conjecture that an investor with more variable performance is more likely to stop trading. Again, there is no reason to believe that the consistency of an investor's past performance should directly affect *changes* in his future performance or disposition effect.

We construct the sample to be used in the selection model as follows. An account observation is added to the selection sample if it places one or more trades in a given year. This differs from our main sample, where we require investors to have placed at least seven round-trip trades in order to estimate either the average performance or the disposition coefficient. Once an account is added, it remains in the selection sample until 2003, which is the end of our data. In some years, an account will have placed enough round-trip trades to be included in our hazard regressions, so the data will include a performance average and disposition estimate for this account. However, each year we will also have data on many accounts for which we do not have performance and disposition estimates. If estimates are available, we treat the account as having been selected into our data.

2.1.4 Predictions about heterogeneity and survival

In order to differentiate between H2 and H3, we examine the role of investor heterogeneity and survival in our data. For these to be the important, the heterogeneity for which we want to control must be at least partially observable. This observation leads to our first prediction about heterogeneity:

Prediction 3. *There is substantial predictable heterogeneity in investors' performance and behavioral bias. Some heterogeneity is correlated with experience.*

We test this conjecture by sorting investors into subsamples based on various characteristics that are ex-ante likely to be related to their financial sophistication. For example, we sort investors by their wealth (proxied by each investor's average daily portfolio value), by whether or not they trade options, and by several other characteristics.[7] We then estimate the average disposition and performance of each subsample, testing whether there is a significant difference between the ex-ante sophisticated and unsophisticated subsamples. We also use each subsample to estimate our simple learning regressions, and we look at the experience coefficients in these subsample regressions to test whether learning across groups occurs at the same rate.

Our next prediction examines the effect on learning of individual heterogeneity more directly:

Prediction 4. *Accounting for individual heterogeneity explains a significant portion of the learning implied by simple models.*

We examine this prediction by estimating our learning regressions once again, this time controlling for investor heterogeneity by including individual fixed effects. We also include year fixed effects to adjust for time-series variation in average market returns. We compare the coefficients of our fixed effects regression to the simple model that we previously estimated.

We go further, exploring our data for potentially important survivorship effects. For survivorship effects to be important, there must be significant attrition in the data, which we confirm by testing our next prediction:

[7] Our proxy for wealth includes only the equity holdings of investors. This measure could be negatively correlated with risk aversion since for any level of total wealth an investor who allocates more money to equities may be more risk-averse. However, even if this is the case, cross-sectional variation in risk aversion cannot explain changes in the disposition effect, performance and survival rates over time.

Prediction 5. *Many new investors will cease trading within a short period of time. Those who remain will trade more over time.*

We examine this prediction by looking at the rate at which investors who are in our sample in one year (having placed seven or more round-trip trades that year) continue to be in the sample in subsequent years. The prediction that surviving traders will trade with more intensity is a feature of standard learning models like that of Mahani and Bernhardt (2007). In this model, investors do not initially know their type. As investors trade, they update their subjective probabilities of being skilled, learning about their inherent ability. Investors with a sufficiently low probability of being skilled eventually cease trading while those that survive increase the intensity of their trades. We examine the intensity part of this prediction by looking at the typical pattern of trading volume for a new account that continues to trade.

Finally, we examine how investor survivorship affects our learning estimates in our last prediction:

Prediction 6. *Accounting for survivorship in addition to individual heterogeneity explains a significant portion of the learning implied by simple models.*

We carefully control for survivorship bias by estimating our learning regressions with the modified Heckman (1976) correction. It is important for us to control for survivorship bias, since it is clear that investors with weaker performance will be less likely to continue trading long enough for us to estimate their performance in future periods. Comparing the results of our simple models to those that correct for survivorship allows us to estimate what fraction of learning by trading is driven by learning about inherent ability (H2) and what fraction is driven by improved ability (H3).

2.2 Data

The data used here are discussed in Chapter I. In addition to those variables, we use the data to create proxies for wealth and measures of investor sophistication. To construct a proxy for wealth, we use opening balances and subsequent trades to reconstruct the total portfolio holdings of each account on a daily basis. Using these holdings, we approximate wealth as the average daily marked-to-market portfolio value for each investor. We also calculate the average value of trades placed by an investor each year. To measure sophistication, we note that investors who trade options are likely to be more familiar with financial markets. This is particularly true in our setting because many of the options in our data are granted to corporate executives as part of compensation. Therefore, while we do not include options trades in our estimates of disposition, we use whether an investor ever trades options as a proxy for sophistication. We also count the number of distinct securities traded by an investor over the sample period, and use this as a measure of portfolio diversification.

Table 2.1 provides summary statistics for the new accounts in our dataset. New accounts are accounts that place their first trade in 1995 or a subsequent year. (That is, they have no recorded initial positions.) Panel A includes all new accounts that place at least one trade during our sample period (1995–2003), while Panel B gives results only for those new accounts for which we are able to estimate the disposition coefficient at least once. We only attempt to estimate the disposition coefficient if an individual has placed at least seven round-trip trades in a given year, although even with this restriction the procedure to maximize the likelihood function does not always converge. The last two rows of each panel are indicator variables, taking a value of one if the investor: (a) trades options; or (b) is female; and zero otherwise.

Comparing Panels A and B, it is apparent that the subset of investors for whom disposition coefficients are available is somewhat different than the larger population. By construction, the accounts in Panel B place more trades, but they also have larger portfolios, trade larger amounts of money, trade in a wider selection of securities, and are somewhat younger. As well, investors for whom we can estimate disposition are more likely to trade options (17%) than the overall sample (3%). Since we are only able to estimate disposition for investors who trade with some frequency, this likely results from the fact that investors who trade options are simply more likely to trade in general.

Figure 2.1 shows the number of accounts (including both new and existing accounts) that place one or more trades in each year. There is considerable variation in the number of accounts placing trades over time, from a low of 54,196 accounts in 1995 to a high of 311,013 accounts in 2000. Additions of new accounts follows a similar pattern. We discuss entry and exit from the sample in more detail in the next section.

2.3 Results

We present our empirical findings in this section. We begin by presenting the results of our tests relating to performance in Section 2.3.1 and to the disposition effect in Section 2.3.2. Section 2.3.3 examines whether there is heterogeneity in learning and whether our tests that adjust for heterogeneity change our learning estimates. Section 2.3.4 examines the importance and magnitude of survivorship effects, and a number of additional tests are presented in Section 2.3.5.

2.3.1 Performance tests and results

We start by performing a few tests on trader performance to be sure that our performance measure has similar properties to the measures used in the literature. Previous papers, in particular Odean (1998), have shown that average investor performance is worse than that of the market portfolio. Poor performance by average investors is also a prediction of the model in Mahani and Bernhardt (2007). This motivates our first test—on average, individuals do not earn returns in excess of the market return. We test this hypothesis by calculating the average return to a stock purchased by an individual investor net of the market return. Calculating this average at a 30-day horizon (using our convention of using a shorter holding period if the individual sells the stock before 30 days) yields an average return net of the market of -4.9 percent. At a 60-day horizon, the average net return is -10.1 percent. At both of these horizons, returns net of the market return are quite statistically significantly negative.

While the average performance of individual investors is likely to be quite poor, Coval, Hirshleifer, and Shumway (2005) shows that some individuals persistently outperform others. Again, performance persistence among individuals is an implication of Mahani and Bernhardt (2007). We test whether there is any persistence in investor performance in three related ways. For the first approach, we regress each investor's average 30-day return in year t on the investor's average return in year $t-1$ and year fixed effects. Using year fixed effects adjusts for time series variation in average market returns. The estimated coefficient in this regression is 0.183 ($p < 0.0001$), very statistically and economically significant. Our second approach is to calculate each investor's average return in two disjoint time periods, 1995–1999 and 2000–2003. We then calculate the Spearman rank correlation between the return series from the first period with that from the second period. This correlation is 0.164 ($p < 0.0001$), again quite statistically and economically significant. Our third test method involves sorting investors in each year into performance quartiles, and then plotting the average performance of each of those quartiles for the next several years. This plot, which appears in Figure 2.2, again gives evidence that the most successful investors in the past continue to outperform the least successful investors for at least a couple of years. Results calculated with alphas instead of raw returns are qualitatively the same. These results confirm that there is a degree of persistence in individual returns.

Our main prediction on performance is that more experienced investors have better investment performance. For now, we test this prediction by simply regressing performance on experience, experience squared (to allow learning to slow down over time) and some control variables, ignoring any individual heterogeneity or survivorship bias. Columns 1 and 2 of Table 2.3 report the results of this regression. When experience is measured either in number of years or cumulative trades, it is positively and significantly related to average returns. An additional year of experience increases average 30-day post-purchase returns by $41 - 4 = 37$ bp, or approximately 3 percent at an annualized rate. An additional 100 trades increases returns at slightly over one-fourth of this rate. Again, results estimated with alpha instead of raw returns are quite similar (unreported). While these estimates are encouraging, the speed of learning they imply seems almost implausibly large. Taking the regression parameters at face value, an investor with 8 years of experience should outperform a new investor by about 22 percent per year. While we observe some heterogeneity in investor ability (or some performance persistence) it is not nearly large enough to justify these large coefficients.

2.3.2 Disposition tests and results

The disposition effect is quite large in our data. To give an idea of the economic significance of the effect, we present some aggregate evidence of the effect in our data. Figure 2.3 is a plot of the relation between the propensity to sell an existing position (the hazard ratio) and the position's holding period return. To generate this plot, we group all investors and estimate one hazard model each year. We group the data for this procedure so we can estimate a model with many covariates, but almost all of the tests that follow are based on individual-level results. Rather than using only one indicator variable as in Equation (2.3), we use 20 dummy variables corresponding to different 1 percent return 'bins' in this model. In Figure 2.3 we plot the dummy variable coefficients by year. The sum of these coefficients times their corresponding dummy variables is multiplied by the baseline hazard rate to give the actual conditional hazard rate. The conditional hazard ratio is remarkably similar across years. The plot shows an obvious kink in the hazard ratio near zero: investors are clearly more likely to sell a stock if it has increased in value since the purchase date. This provides strong support for the presence of a disposition effect in aggregate, consistent with the extensive literature cited above.

Turning to our main individual-level disposition regressions, we require that an investor place at least seven round-trip trades in a year to be included in the sample, and we run the regression for each investor-year to generate a separate disposition coefficient whenever possible. While this filter drastically reduces our sample size, it is necessary to ensure that our coefficients of interest are identified.

Table 2.2 summarizes the distribution of our disposition estimates, which we use to investigate several of our predictions. Panel A provides information on all investors for whom we have estimates. There are 18,042 observations in our panel, and the number of observations each year rises considerably in the first half of the sample and then declines somewhat in the latter part of our sample. The median disposition coefficient is 1.04, which is economically quite large. This coefficient implies that the median new investor in our data is $e^{1.04} = 2.8$ times more likely to sell a stock whose price is above its purchase price than a stock that has fallen in value since the time of purchase.

Using the estimated standard errors for each investor, we can classify estimates as significant or not at any given confidence level. The last two columns of Panel A show the proportion of investors who have a significantly positive or negative disposition coefficient at the 10 percent level. Over our entire sample period, 41.8 percent of investors have a disposition coefficient that is statistically greater than zero. Panel B gives summary statistics for the other coefficients in the hazard model. None of the controls is statistically significant in the cross-section.

Before we can consider whether investors learn to avoid the disposition effect, we need to argue that the effect is in fact a behavioral bias. Theoretically, there is no particularly well accepted model, either rational or irrational, that produces the disposition effect. It is difficult to imagine a rational model that can produce the effect. It is not particularly difficult to think of a model in which the probability of selling a stock increases in the stock's unrealized return, but it is difficult to think of a model in which the probability of selling rises dramatically as soon as the return becomes positive—that is, a model that predicts a 'kink' at zero, as show in Figure 2.3.

It is also difficult to design a model in which traders with the disposition effect have the characteristics that our traders have. In particular, one necessary condition for disposition to be a behavioral bias is that investors with more disposition have inferior investment performance. If disposition is unrelated to investment performance, investors with the effect would have little incentive to learn to avoid it. To get a sense of how returns vary with disposition, we examine average investor returns across quintiles of the disposition coefficient. In this sort, the disposition coefficients are always estimated one year before the average returns are calculated, so disposition coefficients and average returns are not mechanically correlated in any way. For each quintile, Figure 2.4 graphs the average return earned by investors over different horizons from the purchase date. Returns are substantially higher in the lowest disposition quintile than in the highest disposition quintile. For example, in the 30 days following a purchase, a stock's price increases 46 bp on average when bought by an investor in the lowest disposition quintile, compared to a decline of 54 bp if purchased by an investor in the highest disposition quintile. The differences between high- and low-quintile average returns range from 17 bp at the 10-day horizon to 131 bp at the 45-day horizon. These differences are both economically and statistically large. These results are consistent with the claim that individuals with high disposition effect coefficients have relatively poor investment performance. They are also consistent with the disposition effect being a behavioral bias that investors want to learn to avoid.

Another necessary condition for disposition to be a behavioral bias is that disposition is a somewhat stable, predictable attribute of a particular investor. We test this conjecture by estimating the disposition effect at the investor level in adjacent time periods. Each set of estimates comes from a completely disjoint dataset. Any trades that are not closed at the end of the first period are considered censored in the model estimated with first period data. Therefore, any trades that are not closed at the end of the first period are completely ignored in the model estimated with second period data. We explore the stability of disposition coefficients by estimating the rank correlation of account-level disposition coefficients over the two periods, testing whether the rank correlation is significantly different from zero. We estimate the rank correlation between an investor's disposition coefficient in year t and their coefficient in year $t-1$ to be 0.364, suggesting that there is a fair degree of persistence in the individual's disposition coefficient. This correlation is extremely statistically significant.

Taken together, Figures 2.3 and 2.4, Table 2.2 and the correlation in the previous paragraph provide strong evidence that the disposition effect is a widespread and economically important behavioral bias that is present in each year of our study.

Finally, we examine whether more experienced investors are more likely to avoid the disposition effect, testing our main prediction related to disposition. Again at this point, we simply regress disposition coefficients on our experience variables, those variables squared, and some control variables. Columns 3 and 4 of Table 2.3 present our results for the disposition learning regressions. To reduce the weight given to disposition coefficients that are not estimated very precisely, we estimate the regressions with weighted least squares, where the weights are proportional to $1/\widehat{\text{Var}(\beta^d)}$ from our hazard regression in Equation (2.3). The base case (Column 3) shows that disposition declines with experience ($\beta_1 < 0$). Moreover, investors tend to slow down in their learning as they gain experience since $\beta_2 > 0$. Frequent traders, investors who trade more securities, and investors who earned higher returns in the previous year all have lower levels of disposition, but even with these controls our base results are qualitatively unchanged.

Column 4 indicates that an additional 100 trades reduces the disposition coefficient by 0.041, which is similar to the coefficient on Experience in Column

3. In other words, a year of experience or 100 trades have approximately the same effect on disposition. In each of the specifications the estimated YearsTraded and CumulTrades coefficients are statistically significant at the 1 percent level. Economically, however, our results suggest that investors learn relatively slowly. Specifically, the estimates in Column 3 suggest that an additional year of experience corresponds to a reduction in the disposition coefficient of approximately 0.05. To provide some context for this estimate, note that the unconditional median disposition coefficient in our sample is 1.04. An extra year of experience decreases this by about five percent.

As discussed in Section 1.2, we also estimate the yearly hazard models presented in Table 2.4, which are comparable to the model of Feng and Seasholes (2005).[8] These estimates are from pooled hazard models estimated each year, in which all individuals are treated as if they were just one person. Experience is interacted with an indicator for whether the price is above the purchase price, and the coefficient on this interaction term is interpreted as a learning coefficient. These models again give evidence of learning. However, the learning coefficient estimates are quite variable over time (0.149 to 0.063), and they are statistically insignificant in three of the nine years. Furthermore, the average disposition coefficient of an investor is estimated in aggregate to be around 0.65 in this model. This suggests that, depending on the period chosen, an additional year of experi-

[8]These regressions are essentially the same as the empirical strategy employed by Feng and Seasholes (2005), except we use a proportional hazard regression rather than the parametric Weibull model. Our indicator variable is the same as their 'Trading Gain Indicator' (TGI). The experience variable we use for this model is the total number of trades placed before the current trade, rather than CumulTrades, which is the total number of trades placed in previous calendar years.

ence corresponds to a reduction in the disposition coefficient of approximately 10 percent to 25 percent, significantly higher than our regression estimate. Table 2.4 also lists the number of observations available each year, and the fraction of the observations that are censored, or the purchased stocks that are not sold by the end of each year.

The low level of the average disposition effect, the high variability in the annual learning estimates, and the high level of learning found both by Feng and Seasholes (2005) and in our implementation of their model suggest that there are significant differences between our approach and theirs. One important difference is that, since we estimate a different disposition coefficient for each individual, those who place a large number of trades have the same weight in our analysis as those who place seven or eight trades. In Feng and Seasholes (2005), an individual's weight is proportional to her trading volume. Another difference is that we estimate learning over a much longer period of time, since Feng and Seasholes (2005) only have about two years of transactions data. Thus, the year-to-year variation in the annual estimates in Table 2.4 is much less of a concern for our analysis.

2.3.3 Heterogeneity in learning

To examine if there is significant predictable heterogeneity among individuals in our sample, we separate investors into a number of different groups by observable characteristics that we believe are related to their financial sophistication. We examine the average disposition and performance of each of these groups, confirming that our priors are correct and demonstrating that there is significant heterogeneity in the data. We also estimate the simple learning regression for disposition and returns for each group, predicting that the less sophisticated investor groups will learn faster than the more sophisticated investor groups. Importantly, we do not classify investors on the basis of the estimated disposition coefficient, β^d, because of concerns about measurement error. That is, the most extreme disposition estimates are likely those with the most error, and we would therefore expect these accounts to see a decrease in disposition in future years, even if these investors are not really learning. To avoid sorting on measurement error, we focus instead on observable variables that are ex-ante related to performance and disposition.

Each row of Table 2.5 displays the mean of the disposition coefficient and the average returns (or performance) of each group, as well as the regression coefficient of these variables on YearsTraded and the number of observations used in the calculations. Results for disposition are shown in Columns 1 and 2, and for returns in Columns 3 and 4. We consider investors ex-ante likely to be relatively sophisticated if they trade options, have significant wealth, are men, or have had relatively good past performance.[9] Looking at the table, it is clear that the means for each of our sophistication subgroups is significantly different, and each change in the mean across subgroups is of the sign we expect. We also expect that less sophisticated investors will learn faster than more sophisticated investors. In each pair of rows of the table this prediction is confirmed. In most cases there is a clear

[9] We classify investors who make excess profits that are in the top quarter of the entire market in the first two years of their trading as 'winners.' Our results are not sensitive to alternative definitions of winners, such as using a one- or three-year classification period, or above-median excess returns.

difference between the unsophisticated investors, who learn to avoid the disposition effect at a rate of about 10 percent per year, and sophisticated investors, for whom the learning coefficient is often insignificant.

It is also plausible that investors learn more when the market in general is not doing well. During periods of high market returns, investors' incentives to learn about their biases could be reduced if they attribute their success to their ability, similar to the behavior modeled in Zingales and Dyck (2002) in the context of media and bubbles. Thus, we should find that investors are more likely to learn when the markets are not doing well rather than when they are. To test this we define the state of the market as an 'up-market' if the excess return on a broad Finnish index is positive in a given year, and as a 'down-market' if the excess return is negative. In the last two rows of the table we re-estimate our base regressions for each of the two states of the market and find that, as we hypothesized, individuals learn to avoid the disposition effect primarily when the market is not doing well.

These results are consistent with Prediction 3 and suggest that there is substantial heterogeneity both in initial ability (performance and disposition) and in rates of learning. It is unsophisticated investors and investors who start out with poor returns who learn most. These results suggest that considering heterogeneity in our learning estimates will be important.

We control for time-invariant unobserved individual characteristics by including individual fixed effects in our learning regressions. We also control for market returns and any other time-varying features of performance or disposition by including year fixed effects in the regressions. Specifically, we estimate:

$$y_{i,t+1} = \alpha_i + \beta_1 \text{Experience}_{i,t} + \beta_2 \text{Experience}^2_{i,t} + \delta X_{i,t} + \gamma_t + \epsilon_{i,t}, \tag{2.5}$$

where α_i is the individual specific fixed effect, γ_t is the time fixed effect and $X_{i,t}$ are other controls. The performance results, reported in Columns 1 and 2 of Table 2.6, suggest that an investor with one year of experience will earn 22 bp more than an inexperienced investor over a 30-day horizon. Column 2 indicates that a similar increase in returns comes from an additional 400 trades. The disposition results appear in Columns 3 and 4. While YearsTraded is no longer significant in these regressions, Column 4 suggests that 100 trades reduces the disposition coefficient by approximately 0.03. Comparing these estimates with those reported in Table 2.3, we find that the learning estimates after controlling for individual fixed effects fall roughly by one-half. Consistent with Prediction 4, this suggests that though investor performance improves and disposition declines with experience, accounting for individual heterogeneity reduces the estimates by about 50 percent.

2.3.4 Survivorship effects

In this section, we explore our data for potentially important survivorship effects. For these effects to be important, there must be significant attrition in the data. In other words, many new investors should cease trading within a short period of time. The model of Mahani and Bernhardt (2007), which suggests that attrition is driven by investors learning about their inherent ability, also predicts that those traders who remain will trade more intensely over time.

We first present some overall attrition evidence by examining the rate at which investors who are in our sample in one year (having placed 7 round-trip trades in one year) fail to place any trades during the rest of our sample period. Since the rest of the sample period changes from year to year, the earlier years of our sample period provide more reliable estimates of true exit rates than the later years. Figure 2.5 shows that attrition is a significant feature of our data. Approximately 25 percent of those traders who enter the sample in one year fail to ever trade again. Of traders who trade for two or three years, about 5 percent permanently exit the sample. This is consistent with our Prediction 5.

To check whether trading intensity changes over time, we estimate regressions with both the number of trades placed and total trade value as dependent variables, and experience and year dummy variables as explanatory variables. Including year dummies ensures that market-wide changes in stock characteristics will not contaminate our results. We plot the results of these regressions in Figure 2.6. The plot clearly shows that, conditional on survival, trading intensity increases over time. This supports the part of Prediction 5 that has to do with trading intensity, which is consistent with the model of Mahani and Bernhardt (2007).

In some of our last sets of results, we carefully account for a survivorship effect by estimating our learning regressions using the method proposed by Wooldridge (1995), which is a standard Heckman (1976) correction modified to account for individual time invariant heterogeneity. Results from the selection model, with two-step efficient estimates of the parameters and standard errors, are given in Table 2.7. The first-stage selection model uses 30,218 observations, while the second-stage regression (in first differences) use only 6,511 observations in the performance regression and 8,818 observations in the disposition regressions.

We estimate the first-stage regression for each year and construct inverse Mills ratios for each year. For brevity, we only report one set of pooled first-stage estimates in Column 1 of Table 2.7. Results for each of the years are qualitatively similar to those reported. We find strong evidence that as investors get bad returns they cease trading. In particular the estimate on $\bar{R}_{t-1}$ is positive and significant. The estimate is also economically meaningful and suggests that, keeping other explanatory variables at their mean levels, a decrease in returns of one standard deviation increases the probability that the individual will cease to trade next period by around 15 percent. This is strong evidence for H2. As low ability investors trade, they learn about their inherent ability and cease trading. More successful investors continue to trade actively.

The other coefficient estimates reported in the first column of Table 2.7 also seem sensible: investors are more likely to remain in the sample and trade if they hold relatively diversified portfolios and have relatively more trading experience. Importantly, coefficient estimates for both of our instruments are statistically significant and of the predicted sign. Specifically, inheriting shares increases the probability that the individual will continue trading the next period and higher variability in past performance increases the probability that the individual ceases to trade. Both the instruments also have an economically significant impact. For instance, keeping other variables at mean levels, inheriting shares increases the probability that the investor will continue trading in the next period by 5 percent.

In line with Prediction 6, we find that accounting for selection has a significant impact on our learning estimates. Column 2 uses performance as the dependent variable in a regression of the form of Equation (2.4). Comparing the estimates in Table 2.7 to the simple model reported in Table 2.3, the coefficient on YearsTraded is no longer statistically significant, and the coefficient on cumulative trades is reduced by about 90 percent.[10] When we use disposition as the dependent variable in Column 3, the coefficient is reduced by slightly more than 50 percent. The joint tests of statistical significance of the inverse Mills ratios in Columns 3 and 4 also show that accounting for sample selection is important. Unreported results in which YearsTraded and CumulTrades are included in separate regression models yield almost the same coefficients, but regressions that include both variables are reported for brevity.

Our results suggest that accounting for selection is important and significantly affects inferences about learning. Investor heterogeneity and survivorship effects account for something on the order of one-half to three-quarters of the learning estimates found in simple and aggregate models. This translates directly into slower

[10] The coefficients on YearsTraded in these regressions must be viewed with caution. Because the change in YearsTraded is always equal to exactly one year, the coefficient on YearsTraded can only be identified if we leave the fixed effect for one year out of the regression. We leave the fixed effect for 1997 out of the regression, so the coefficient we report can be thought of as the learning coefficient for that year. For robustness, we also estimate standard Heckman selection models without either individual or year fixed effects. The coefficients on YearsTraded in these models are 0.31 in the performance regression and -0.019 in the disposition regression. Both of these coefficients are marginally statistically significant (at 10 percent). The magnitude and significance of these coefficients is consistent with our other results.

learning than that inferred from simpler models. Taking the disposition-learning coefficient, for example, 100 trades corresponds to an improvement of about 0.04 in the simplest model, an improvement of about 0.03 in the model with individual fixed effects, and an improvement of about 0.02 in the survivorship/fixed effects model. Roughly speaking, if it takes about 100 trades to improve about 4 percent in the simple model, it takes about 200 trades to achieve the same improvement after adjusting for survivorship and individual heterogeneity.

Our estimates suggest that the fraction of learning that is driven by investors learning about their inherent ability (i.e., H2) by trading is large. After adjusting for this type of learning, the portion of learning that is due to investors learning to improve their ability over time (i.e., H3) is significantly different from zero, but not excessively large. In other words, there is support for H3, though accounting for H2 significantly reduces estimates of how quickly investors become better at trading. Overall, our findings are consistent with the three hypotheses that were outlined in Section 2.1.

2.3.5 Other Tests

In this section we conduct some additional tests related to our main predictions. First, it is possible that the performance improvement with experience that we find is not entirely due to learning of the types we considered but rather due to a change in the risk preferences of investors over time. To address this possibility, we estimate our learning regressions with risk-adjusted returns (or 30-day alphas) instead of raw returns, and we regress the average factor betas of stocks purchased by investors on experience and our control variables. Our regressions also control for survivorship and individual and year fixed effects. The results of our regressions appear in Table 2.8. The table clearly shows that risk-adjusted returns improve with experience, with a coefficient that is actually larger than the coefficient we estimate for raw returns. Looking at the coefficients on average factor betas makes it clear why this is the case. With more experience, investors are actually both improving raw returns and taking less risk, or purchasing stocks with lower factor betas. This result is particularly strong for the market (RMRF) and size (SMB) factor betas. Thus, it appears very unlikely that our raw return learning results are driven by changes in risk preferences with experience.

Second, in unreported results, we substitute the market return for each stock's return to see if individuals learn to time the market. If investors are learning to identify good times to buy then the market as a whole will tend to increase after their purchases; if instead they are learning to select stocks, we will not find evidence of learning when we look only at market returns. In fact, we find that the coefficient estimates on experience variables are insignificant, which suggests that performance improved because investors became better at stock selection. Third, we also conduct the tests on survivorship using the Wooldridge (1995) method, taking data on investors who resume trading after ceasing to trade for a few years (the tests reported in the last section had dropped such investors, using only observations in two consecutive years). Including these investors increases the sample by around 250 observations in the second stage but does not affect the nature of the results reported. Fourth, all of the results on disposition remain qualitatively unchanged if we include a 'December dummy' in (2.3) or remove partial sales from our sample. This rules out tax-motivated selling or rebalancing as possible explanations for the disposition effect. Finally, in all the fixed effect regressions that control for individual heterogeneity, we cluster the standard errors at the individual level and find that our results are unaffected.

2.4 Conclusion

We examine learning in a large sample of individual investors in Finland during the period 1995–2003. We correlate performance and disposition with investor experience and investor survival rates to determine whether and how investors learn by trading. We find that performance improves and the disposition effect declines as investors become more experienced, suggesting that investors learn by trading. We differentiate between investors learning about their inherent ability by trading and learning to improve their ability over time by accounting for investor attrition. We find that a substantial part of this learning occurs when investors stop trading after learning about their inherent ability rather than continuing to trade and improving their ability over time. By not accounting for investor attrition and heterogeneity, the previous literature significantly overestimates how quickly investors become better at trading.

Our results suggest a number of interesting implications. First, since investors who continue trading learn slowly and there is great deal of turnover in the investor population, it is likely that behavioral biases are an important feature of financial markets. Agents do not learn fast enough to make it impossible for biases to affect asset prices. Second, while it would be wonderful to know how quickly investors who cease to trade would learn if they chose to continue trading, we have no way to estimate this speed. If we assume that those who continue trading learn more quickly than those who cease to trade, policy makers might enhance welfare by devising screening mechanisms, or tests that measure and reveal inherent investing ability. Allowing unskilled investors to learn of their poor ability without incurring significant costs might be more valuable than encouraging people to become active investors. Third, an open question in the literature is why there is such high trading volume, particularly among seemingly uninformed individual investors. Our results indicate that such trading may be rational; investors may be aware that they will learn from experience, and choose to trade in order to learn. Our results also suggest that differences in the expected performance of investors may arise from different experience levels. Finally, if many inexperienced investors begin trading around the same time, their trades could lead to time-varying market efficiency. Our evidence is therefore consistent with the recent results of Greenwood and Nagel (2007), and the more general discussion found in Chancellor (2000) and Shiller (2005).

Table 2.1: Summary Statistics

This table presents summary statistics for our data. Panel A includes all individual accounts in our data that started trading during the sample period. Panel B gives results just for those accounts for which we are able to estimate at least one disposition coefficient. We only estimate the disposition coefficient if an individual has placed at least seven round-trip trades in a given year. Number of trades is the total number of trades placed by an investor during the sample period. Average portfolio value is the average marked-to-market value of an investor's portfolio using daily closing prices.

Panel A: Entire Sample (322,454 accounts)

	Mean	25th Pctl	Median	75th Pctl
Number of years with trades	1.9	1.0	1.0	2.0
Number of securities traded	3.5	1.0	1.0	3.0
Number of Trades	15.4	1.0	3.0	8.0
Average value of shares traded, EUR	3,447	808	1,653	3,310
Average portfolio value, EUR	11,588	1,470	2,794	5,856
Age in 1995	39.3	27.0	39.0	51.0
Gender (1=female)	0.39			
Trades options (1=yes)	0.03			

Panel B: Accounts with Disposition Estimates (11,979 accounts)

	Mean	25th Pctl	Median	75th Pctl
Number of years with trades	4.4	3.0	4.0	6.0
Number of securities traded	22.3	12.0	18.0	28.0
Number of trades	222.3	68.0	117.0	224.0
Average value of shares traded, EUR	5356	1855	3235	5759
Average portfolio value, EUR	58828	5102	11483	26147
Age in 1995	35.3	27.0	34.0	44.0
Gender (1=female)	0.15			
Trades options (1=yes)	0.17			

Table 2.2: Disposition Estimates

This table reports a number of summary statistics for our estimates of the disposition effect. β^d is the coefficient in the hazard regression,

$$\lambda_{i,j}(t|x_{i,j}(t)) = \phi_i(t)\exp\Big\{\beta_i^d I(R_{i,j}(t) > 0) + \beta_i^r \bar{R}_{M,t} + \beta_i^s \sigma_{M,t} + \beta_i^V V_{M,t}\Big\}.$$

We estimate this model for each account-year with seven or more round-trip trades. Panel A reports on the cross-section of coefficients estimated each year. The columns labeled 'positive' and 'negative' report the proportion of investors with a statistically significant coefficient, where significance is measured at the 10% level using standard errors obtained from the maximum likelihood estimation of the hazard model. Panel B provides information for all the variables in the hazard model, where β^r, β^s, and β^V are the coefficients on 5-day moving averages of market returns, market returns squared, and market volume, respectively.

Panel A: All Accounts with Disposition Estimates

		β^d estimate			Significant at 10%	
Year	N Obs	10th Pctl	Median	90th Pctl	Positive	Negative
1995	25	-0.34	0.72	2.37	28.0%	4.0%
1996	89	-0.38	0.99	2.39	38.2%	2.2%
1997	248	-0.91	0.96	2.31	33.1%	3.2%
1998	695	-0.69	1.05	2.66	35.4%	2.6%
1999	1958	-0.51	1.08	2.57	40.4%	1.7%
2000	5961	-0.37	1.00	2.54	41.5%	1.5%
2001	3732	-0.29	1.01	2.53	42.7%	1.3%
2002	2649	-0.31	1.10	2.70	45.6%	1.1%
2003	2685	-0.33	1.09	2.53	41.6%	1.3%
All years	18042	-0.36	1.04	2.57	41.8%	1.5%

Panel B: Hazard Function Estimates

Variable	Mean	t-stat	10th Pctl	Median	90th Pctl
β^r	-0.04	-0.27	-0.65	0.03	0.93
β^s	0.01	0.46	-0.05	0.00	0.04
β^V	0.25	9.14	-2.62	0.24	3.08
β^d	1.32	6.81	-0.36	1.04	2.57

Table 2.3: Simple Learning Model: Estimates at Individual Level

This table presents results for regressions of the form

$$y_{i,t+1} = \alpha + \beta_1 \text{Experience}_{i,t} + \beta_2 \text{Experience}^2 + \delta X_{i,t} + \epsilon_{i,t}$$

where the dependent variable is either the investor's average 30-day return following purchases ($\bar{R}_{i,t+1}$) or the investor's disposition coefficient ($\beta^d_{i,t+1}$). Experience is measured by either years of experience (YearsTraded) or cumulative number of trades placed (CumulTrades). $X_{i,t}$ is a vector of controls including the number of trades placed by the individual in a given year (NumTrades), the number of securities held by the individual in a given year (NumSec), and the individual's average total daily portfolio value (PortVal). Data are from the period 1995 to 2003. Standard errors are in parentheses, and ***, ** and * denote significance at 1%, 5% and 10%, respectively.

	Dependent Variable			
	$\bar{R}_{i,t+1}$		$\beta^d_{i,t+1}$	
	(1)	(2)	(3)	(4)
YearsTraded_t	0.414 (0.160)***		-0.050 (0.014)***	
YearsTraded^2_t	-0.043 (0.027)		0.007 (0.002)***	
CumulTrades_t ($\div 10^2$)		0.110 (0.043)***		-0.041 (0.003)***
CumulTrades^2_t ($\div 10^4$)		-0.0005 (0.0003)*		0.0002 (0.00002)***
NumSec_t	0.096 (0.012)***	0.104 (0.012)***	-0.011 (0.001)***	-0.012 (0.001)***
NumTrades_t	-0.031 (0.005)***	-0.043 (0.007)***	0.0006 (0.0005)	0.005 (0.0006)***
PortVal_t ($\div 10^6$)	0.104 (0.061)*	0.099 (0.058)*	-0.006 (0.003)**	-0.003 (0.003)
$\bar{R}_{t-1}$			-0.005 (0.001)***	-0.005 (0.001)***
Observations	13404	13404	17715	17715
R^2 (%)	1.6	1.7	1.1	1.6

Table 2.4: Simple Learning Model: Disposition Estimates at Aggregate Level

This table presents learning estimates from pooled proportional hazards models, using a method similar to that of Feng and Seasholes (2005). Each year we pool the trades of all investors, treating them as if they were just one individual and estimating one learning coefficient for the entire population. The model is,

$$\lambda(t) = \phi(t)\exp\{\beta_1 I(R_{i,j}(t) > 0) + \beta_2 \text{Exper} + \beta_3 \text{Exper}^2 + \beta_4 \text{Exper}[I(R_{i,j}(t) > 0)] + \beta_5 \text{Exper}^2[I(R_{i,j}(t) > 0)]\},$$

where Exper is measured in years since first placing a trade and $I(R_{i,j}(t) > 0)$ is an indicator variable that takes a value of one when a stock has increased in price since its purchase date. For brevity, only the β_4 coefficient estimates are reported. Estimated β_5 coefficients are all insignificant. We also report the number of trades, or observations, considered by the model (in thousands of trades), the percentage of observations censored (trades not closed by the end of the year), and the number of accounts contributing observations to the model. Standard errors are in parentheses, and ***, ** and * denote significance at 1%, 5% and 10%, respectively.

Year	1995	1996	1997	1998	1999	2000	2001	2002	2003
β_4	-0.074	-0.093	-0.096	-0.149	-0.031	-0.013	-0.063	-0.069	-0.123
	(0.039)*	(0.062)	(0.031)***	(0.022)***	(0.024)	(0.011)	(0.011)***	(0.012)***	(0.012)***
Trades	6.6	14.5	26.6	44.6	99.6	251.9	161.9	115.6	131.8
Censored	16%	19%	19%	19%	17%	16%	17%	18%	20%
Accounts	384	713	1360	2412	4953	11585	7445	5416	6056

Table 2.5: Heterogeneity in Learning

This table reports both means and simple learning coefficient estimates from regressions of the form,

$$y_{i,t+1} = \alpha + \beta_1 \text{YearsTraded}_{i,t} + \beta_2 \text{YearsTraded}^2_{i,t} + \delta X_{i,t} + \gamma_t + \epsilon_{i,t},$$

conditioned on a number of variables that ex-ante might be correlated with trader sophistication. The variable of interest is either the disposition coefficient ($\beta^d_{i,t+1}$) or returns ($\bar{R}_{i,t+1}$). For brevity we only report β_1 coefficients in the table. We classify investors as 'Trades options' if they trade in options at any point during our sample. Similarly, investors are classified as 'wealthy' if they are in the top 25th percentile of average portfolio value. We define the state of the market as 'up' if the excess return on a broad market index is positive, and 'down' if it is negative. Finally, we classify investors who make excess profits that are at or above the 75th percentile of all investors in the first two years of their trading as 'winners.' We include year dummies in all the regressions. Data are from the period 1995 to 2003. Standard errors are in parentheses and ***, ** and * denote significance at 1%, 5% and 10% respectively. All group means are significantly different at 1% level.

	Dependent Variable				
	$\beta^d_{i,t+1}$		$\bar{R}_{i,t+1}$		
Classification	Mean	YearsTraded	Mean	YearsTraded	N Obs
No options trades	1.17	-0.096	-0.81	0.501	26698
		(0.024)***		(0.184)***	
Trades options	0.99	-0.031	0.10	0.233	8211
		(0.032)		(0.299)	
Not wealthy	1.14	-0.101	-1.18	0.389	26176
		(0.023)***		(0.169)**	
Wealthy	1.11	-0.015	1.08	0.191	8733
		(0.043)		(0.339)	
Females	1.24	-0.140	-0.38	0.482	4703
		(0.047)***		(0.129)***	
Males	1.11	-0.078	-0.62	0.361	30206
		(0.020)***		(0.169)**	
Not winners	1.15	-0.840	-0.62	0.732	32456
		(0.021)***		(0.161)***	
Winners	0.96	-0.039	-0.03	-0.641	2453
		(0.074)		(0.703)	
Down market	1.14	-0.120	-0.70	0.619	27074
		(0.022)***		(0.179)***	
Up market	1.08	0.028	1.77	0.114	7835
		(0.077)		(0.471)	

Table 2.6: Learning with Individual Fixed Effects

This table reports estimates of regressions of the form

$$y_{i,t+1} = \alpha_i + \beta_1 \text{Experience}_{i,t} + \beta_2 \text{Experience}^2_{i,t} + \delta X_{i,t} + \gamma_t + \epsilon_{i,t},$$

where Experience is measured by either years of experience (YearsTraded) or cumulative number of trades placed (CumulTrades). The dependent variable in models (1) and (2) is individual i's average return in the following year, $\bar{R}_{i,t+1}$, and in models (3) and (4) it is the individual's disposition coefficient, $\beta^d_{i,t+1}$. $X_{i,t}$ is a vector of controls including the number of trades placed by the individual in a given year (NumTrades), the number of securities held by the individual in a given year (NumSec), and the individual's average total daily portfolio value (PortVal). We also include year dummies and individual-specific intercepts in each regression. Data are from the period 1995 to 2003. Standard errors are in parentheses and ***, ** and * denote significance at 1%, 5% and 10% respectively.

	Dependent Variable			
	$\bar{R}_{i,t+1}$		$\beta^d_{i,t+1}$	
	(1)	(2)	(3)	(4)
YearsTraded_t	0.249 (0.401)		-0.010 (0.038)	
YearsTraded^2_t	-0.022 (0.038)		0.010 (0.003)***	
CumulTrades_t ($\div 10^2$)		0.058 (0.059)***		-0.030 (0.005)***
CumulTrades^2_t ($\div 10^4$)		-0.0009 (0.0004)**		0.0002 (0.00003)***
NumSec_t	0.003 (0.024)	0.002 (0.024)	-0.008 (0.002)***	-0.008 (0.002)***
NumTrades_t	-0.021 (0.010)**	-0.029 (0.010)***	-0.0007 (0.0009)	0.0003 (0.0009)
PortVal_t ($\div 10^6$)	0.618 (0.602)	0.587 (0.601)	-0.067 (0.602)	-0.065 (0.601)
$\bar{R}_{t-1}$			-0.002 (0.001)*	-0.002 (0.001)*
Observations	13404	13404	17715	17715
Adjusted R^2 (%)	45.7	45.9	38.7	38.6
Year Fixed Effects	Yes	Yes	Yes	Yes
Individual Fixed Effects	Yes	Yes	Yes	Yes

Table 2.7: Learning with Survival Controls

This table reports estimates of selection model regressions with the fixed effects modification developed by Wooldridge (1995). The regressions are of the form

$$\Delta y_{i,t+1} = \beta \Delta X_{i,t} + \rho_2 d2_t \lambda_{i,t} + \ldots + \rho_T dT_t \lambda_{i,t} + \epsilon_{i,t},$$

where $\lambda_{i,t}$ are the inverse Mills ratios from a year t cross-sectional probit model (the selection model) and $d2_t, \ldots, dT_t$ are time dummies. Including these variables in the learning regression accounts for the impact of the selection equation. The joint test of $\rho_t = 0$ for $t = 2, \ldots, T$ is a test of whether survivorship bias is a concern. The probit model reported in Column 1 is estimated with data from all of the years of the sample pooled together, but the inverse Mills ratios in the second stage estimates are estimated separately each year. The regressions in Columns 2 and 3 are estimated with all the variables in first differences, except the inverse Mills ratios. These first differences add fixed effects to the model. Experience is measured by years of experience (YearsTraded) and cumulative number of trades placed (CumulTrades). The dependent variable in the second stage is either the individual's average return in the following year, $\bar{R}_{i,t+1}$, or the individual's disposition coefficient, $\beta^d_{i,t-1}$. $X_{i,t}$ is a vector of controls including the number of trades placed by the individual in a given year (NumTrades), the number of securities held by the individual in a given year (NumSec), and the individual's average total daily portfolio value (PortVal). $\bar{R}_t$, the investor's average return, is excluded from the return regression due to the econometric problems that arise when lagged dependent variables are included in fixed effects models. Additional variables in the first stage selection model are $I(\text{Inherit} = 1)$, an indicator variable for whether the investor inherited shares in year t, and σ_{R_t}, the standard deviation of the investor's returns. Data are from the period 1995 to 2003. Standard errors are in parentheses, and ***, ** and * denote significance at 1%, 5% and 10% respectively.

	First Stage (All Years)	Second Stage (First Difference)	Second Stage (First Difference)
Dependent variable	$\text{Insample}_{i,t+1}$=1	$\bar{R}_{i,t+1}$	$\beta^d_{i,t+1}$
CumulTrades_t ($\div 10^2$)	0.355	0.010	-0.019
	(0.021)***	(0.004)***	(0.002)***
CumulTrades^2_t ($\div 10^4$)	-0.002	-0.0003	0.0002
	(0.0001)***	(0.0002)	(0.00003)***
YearsTraded_t	0.376	0.894	-0.099
	(0.0196)	(1.766)	(0.204)
YearsTraded^2_t	0.028	-0.015	0.008
	(0.007)**	(0.021)***	(0.006)
NumSec_t	0.187	0.034	-0.010
	(0.001)***	(0.009)***	(0.002)***
NumTrades_t	0.963	-0.100	0.002
	(0.062)***	(0.012)***	(0.003)
PortVal_t ($\div 10^6$)	0.003	0.220	-0.007
	(0.003)	(0.221)	(0.007)
$\bar{R}_t$	0.002		-0.003
	(0.001)***		(0.002)*
$I(\text{Inherit} = 1)$	0.191		
	(0.074)***		
σ_{R_t}	-0.617		
	(0.061)***		
Joint Test of $\rho = 0$		F(8, 6490) = 3.91	F(7, 8799) = 5.29
(from 1996-2003)		Pr>F = 0.0000	Pr>F = 0.0000
Observations	30218	6511	8818
Year Fixed Effects		Yes	Yes
Individual Fixed Effects		Yes	Yes

Table 2.8: Risk Taking and Experience

This table reports the results of fixed effect selection model estimates of regressions of various performance and risk measures on experience measures. The method of these regressions and their associated first-stage estimates are described in Table 2.7. The dependent variables include each investors average 30-day risk adjusted return (alpha), and each investors average beta coefficient on four factors—RMRF, SMB, HML, and UMD. These betas are estimated in the standard way, as described in the text. Each regression includes the control variables and inverse Mills ratios described in Table 2.7, but only the experience variables coefficients are reported in this table. Standard errors are in parentheses, and ***, ** and * denote significance at 1%, 5% and 10% respectively.

	Dependent Variables				
Coefficient	30-day α	β on RMRF	β on SMB	β on HML	β on UMD
YearsTraded	0.58	-0.034	-0.016	0.0544	0.0073
	(0.41)	(0.0108)***	(0.0039)***	(0.0435)	(0.0250)
YearsTraded^2	0.02	0.00246	0.00163	0.0037	-0.00076
	(0.02)	(0.0012)*	(0.0005)***	(0.0028)	(0.0019)
CumulTrades ($\div 10^2$)	0.05	-0.01	-0.002	-0.002	-0.016
	(0.02)***	(0.002)***	(0.0009)***	(0.004)	(.005)***
CumulTrades^2 ($\div 10^4$)	-0.0002	0.00005	0.00001	0.00002	0.00008
	(0.0001)**	(0.00001)***	(0.000006)**	(0.00002)	(0.00002)***
Dep Var Mean	0.07	0.663	0.028	-0.086	-0.054
Standard Dev	3.30	0.288	0.103	0.431	0.440

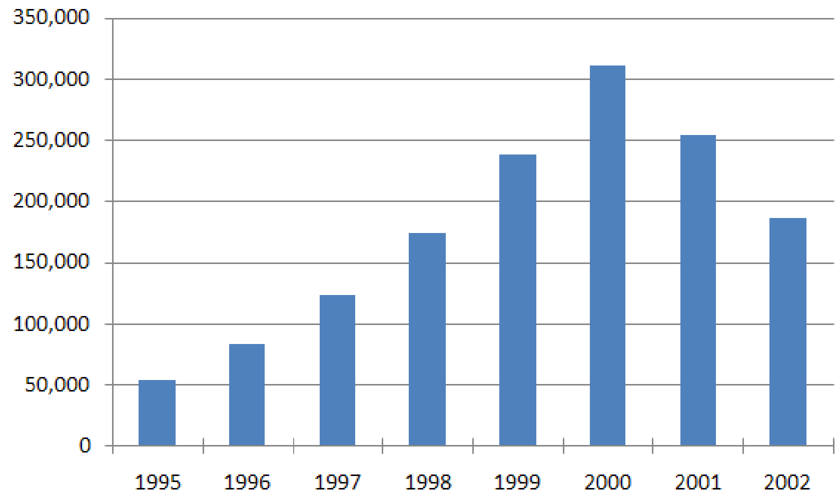

Figure 2.1: Participation By Year

This graph shows the number of accounts that place one or more trades in each year, including both accounts that exist at the beginning of the sample and new accounts. There is considerable variation in the number of accounts placing trades over time, from a low of around 54,196 in 1995 to a high of 311,013 in 2000.

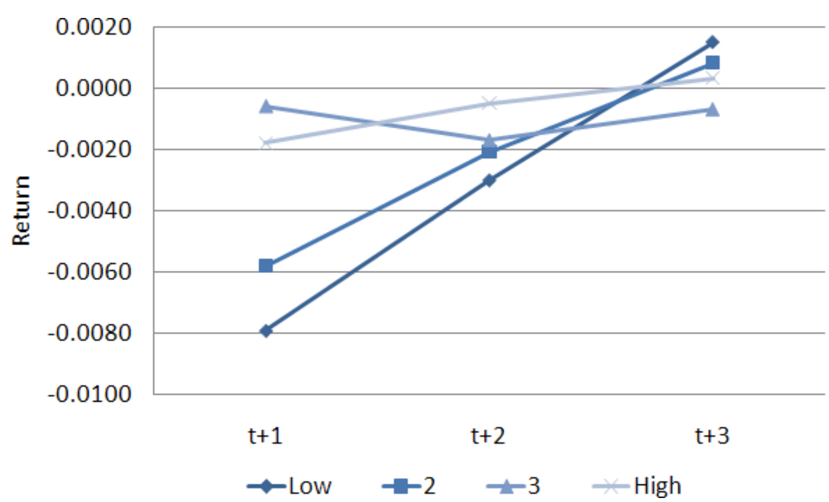

Figure 2.2: Returns Persistence

This figure plots the average 30-day returns earned by investors in years following their first purchase. Investors are grouped into quartiles in their first year of trading, and we then calculate average returns for each quartile in subsequent years. Returns are calculated using the approach discussed in the text, and are demeaned by calendar year, which removes the impact of any year fixed effects. Raw returns are reported here, but the results for risk-adjusted returns are not qualitatively different.

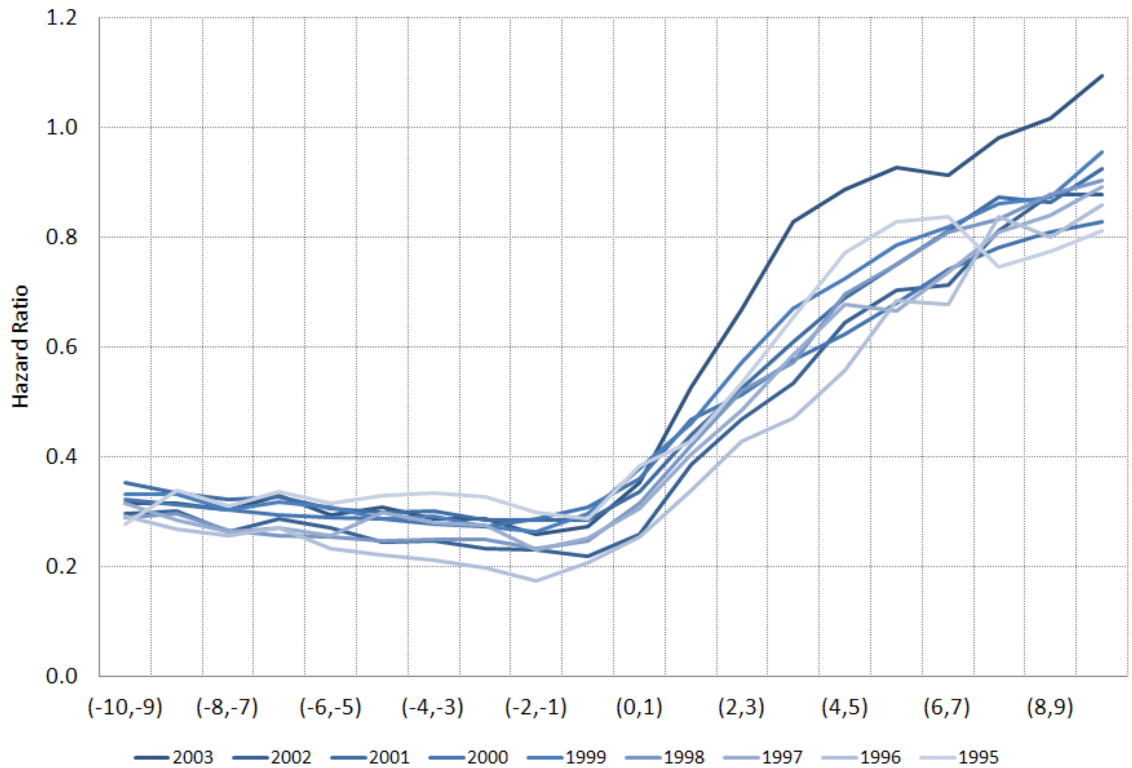

Figure 2.3: The Disposition Effect in Aggregate

This graph shows how the propensity to sell a stock depends on the stock's return since purchase. Each line plots the regression coefficients from one hazard regression modeling the conditional probability of selling a stock. The coefficients correspond to dummy variables for return 'bins' ranging from $[-10, -9)$ percent to $[9, 10)$ percent. In each year, there is a pronounced kink near zero, and the hazard increases rapidly for positive returns.

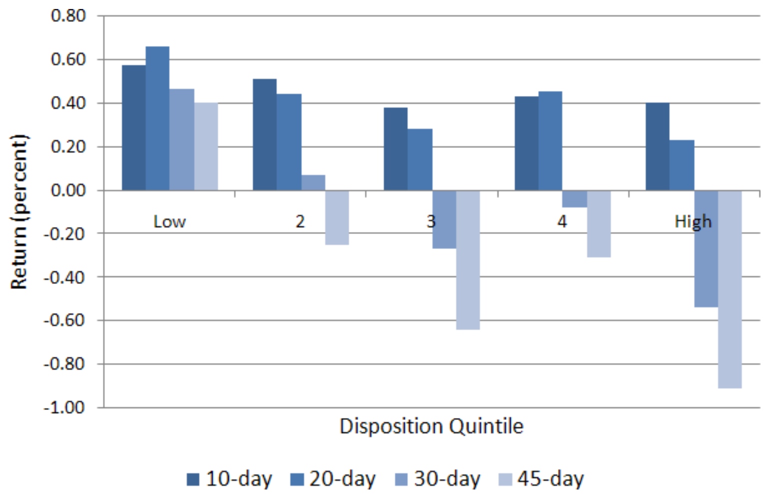

Figure 2.4: Returns by Disposition Quintile

This figure shows average 10-, 20-, 30-, and 45-day returns following a purchase for each disposition quintile. Returns are calculated using the approach discussed in the text. Raw returns are reported here, but the results for risk-adjusted returns are not qualitatively different. Returns earned by the lowest quintile (1) are higher than those earned by the highest quintile (5).

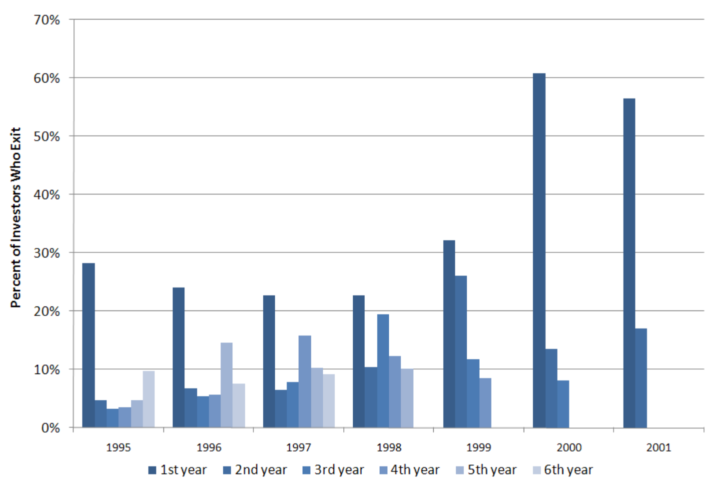

Figure 2.5: Proportion of Accounts Who Exit

This graph shows the percentage of investors who exit in 1 year, 2 years, . . . , or 6 years following the first year in which they place at least seven trades. Exit is defined as placing no further trades during our sample period. Within each year group, the first bar indicates the percentage of investors who exited in the 1st year after placing a trade, the second bar indicates the percentage of investors who exited in the 2nd year, and so on. Data are missing for later years because we do not know how many investors stopped trading after 2003.

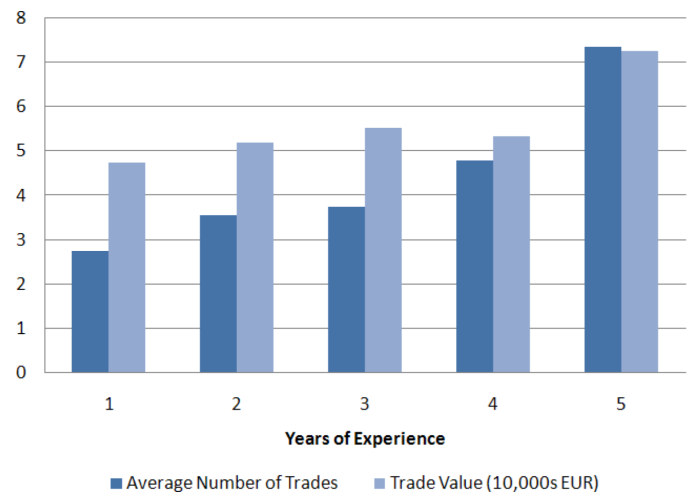

Figure 2.6: Trading Intensity and Experience

This figure shows how trading intensity changes with experience. Intensity is measured as the number of trades placed (dark blue) and the total value of trades placed (light blue; in 10,000's of EUR). Results are demeaned by calendar year to adjust for year fixed-effects. We report results beginning in the investor's first full calendar year in the sample.

Chapter III

Who Trades with Whom?

Models of market microstructure posit the existence three types of financial market participants: informed investors, uninformed "noise traders," and market-makers (Glosten and Milgrom 1985, Kyle 1985). An important question in economics, dating back at least to Keynes (1936) is to what extent prices are distorted by noise traders.[1] Friedman (1953) argues that rational, informed investors quickly exploit arbitrage opportunities caused by mispricing. But De Long, Shleifer, Summers, and Waldmann (1991) show how noise traders can have long-term effects on prices, and Shleifer and Vishny (1997) and Abreu and Brunnermeier (2003) explain why arbitrageurs may be unable to take advantage of known mispricing. If noise traders distort prices, then prices must move in response to their trading. This chapter examines that possibility.

Who are the proverbial noise traders? While they may have an exogenous "liquidity" motive for trade, Black (1986, p. 531) defines noise trading as "trading on noise as if it were information." The literature provides considerable evi-

[1] Keynes (1936, p. 155) describes investing as a "battle of wits to anticipate the basis of conventional valuation a few months hence, rather than the prospective yield of an investment over a long term of years..."

dence that individual investors play this role.[2,3] For example, individual investors make remarkably poor investment decisions. In the data used in this chapter, stocks heavily bought by individuals *underperform* stocks heavily sold by 6.8% over the subsequent year. In sharp contrast, stocks heavily bought by institutions outperform stocks heavily sold by 9.0% (see Table 3.1). There are two possible explanations for these results, which have been previously documented in other data. One, offered by Barber, Odean, and Zhu (2006) and Hvidkjaer (2006), is that individual investors push prices away from fundamentals; the subsequent reversal to fundamental value leads to the documented poor performance. The second, supported by Kaniel, Saar, and Titman (2008) and Campbell, Ramadorai, and Schwartz (2008), is that well-informed institutions buy undervalued stocks from individuals and sell overvalued stocks to individuals, and prices subsequently move toward fundamentals. Under the first explanation, individuals distort prices; under the second, institutional demand for trading is met by individuals whose trading supplies liquidity.[4]

In this chapter, I use the complete daily trading records for all trading in Finland over a nine-year period to examine these competing hypotheses. First, in contrast to the existing literature, I identify how much trading occurs not only *be-*

[2]See Barberis and Thaler (2005), especially Section 7 and the references therein.

[3]Throughout this chapter, I use the terms "individual" and "household" investors interchangeably. In all cases, I am referring to individual traders, and not institutional investors or professional traders.

[4]The first story in particular also requires a considerable amount of herding among individual investors. I provide estimates on the degree of in-group correlation for individual investors below. Several coordinating mechanisms could explain this herding. For example, individuals may notice prices falling and choose to purchase because they think it's a good buying opportunity. This may be thought of as a price-based mechanism. Another example would be retail brokers acting to sell inventory by encouraging clients to buy shares.

tween individuals and institutions, but also *within* each group. I document that institutions are about half as likely to trade with individuals as they are to trade with other institutions. This is particularly interesting given the herding that has been documented among both individuals (Odean 1998) and institutions (Wermers 1999), which would mean that trading between groups is very common—if all institutions are buying, they cannot buy from other institutions. Second, I show that prices move consistently when institutions trade with individuals. In particular, when individuals buy shares from institutions, prices decline; and when they sell shares to institutions, prices increase. Of course, this implies that prices fall when institutions sell shares to individuals, and rise when institutions buy shares from individuals. In other words, institutions move prices and individuals supply liquidity to meet the trading needs of institutions. I confirm these results at daily, weekly, and monthly horizons, and with vector autoregressions. Third, I show that when prices move as a result of trading among individuals, they are more likely to revert than after trading between individuals and institutions or between two institutions. This price reversion is consistent with individuals being uninformed. Moreover, I show that these price reversions occur because institutions subsequently trade with individuals in a direction that moves prices back toward previous levels.

My data include the complete daily trading records of all households, financial institutions and other entities that trade stocks on the Helsinki stock exchange between 1995 and 2003. There are three notable features of these data that make them particularly well-suited to examining the relation between trading and price changes. First, the data include account identifiers that classify the investor as a household, financial institution, nonfinancial corporation, government agency, nonprofit organization, or foreigner. Therefore, there is no need to estimate an investor classification as there is in datasets available for the U.S. Second, whereas data available in the U.S. are either available quarterly or from proprietary datasets covering small samples of traders and/or short time periods, the Finnish data record all transactions placed each day by each investor. This allows me to analyze the interaction of investors at a high frequency without relying on the estimation technique developed by Campbell, Ramadorai, and Schwartz (2008). Third, the data cover a nine-year period for the entire Finnish stock market. The sample includes both the "bubble" period in technology stocks during which many Finnish stocks rose dramatically, as well as periods before and after this rise. This helps ensure that the results are generally applicable to a variety of market conditions, and not driven by rare events.

The poor performance of individual investors documented by Odean (1998, 1999) and Barber and Odean (2000, 2001) can result either because they trade with better-informed institutional investors, or because they push prices above or below fundamentals and subsequently lose money in the ensuing correction. Barber, Odean, and Zhu (2006) and Hvidkjaer (2006) present evidence that the trading of individual investors moves prices, which then slowly revert. Because of constraints on the data available for the U.S. market, these authors adopt a clever strategy to identify the trading of individual investors: they examine the imbalance of buyer- and seller-initiated transactions for small quantities of trades and classify this as the trading of individuals. Barber, Odean, and Zhu show that this order-imbalance is correlated with the order-imbalance among a sample of investors at a discount brokerage firm. A potential problem with this approach is that either individuals or institutions could be initiating these small trades. The identification strategy employed by these authors, however, relies on the assumption that individuals initiate these trades. This could raise concerns about the generality of the results found in these papers. Indeed, in contrast to these results, Kaniel, Saar, and Titman (2008), Campbell, Ramadorai, and Schwartz (2008), and Linnainmaa (2007) all find that individual investors supply liquidity to meet institutional demand for immediacy.[5]

The remainder of the chapter is organized as follows. Section 3.1 develops the hypotheses to be tested in the chapter. In Section 3.2, I motivate my classification method and show how it is implemented. I present my results in Section 3.3. Section 3.4 concludes.

[5]The large literature examining the performance of institutional investors also highlights the importance of using high-frequency data. Increases in institutional ownership are associated with price increases at a quarterly or annual frequency (Nofsinger and Sias 1999, Cai and Zheng 2004). However, it is difficult to tell from quarterly data whether institutional trading within the quarter leads, lags, or is contemporaneous with returns. This makes it difficult to determine causality, particularly since fund managers follow momentum strategies (Carhart 1997, Daniel, Grinblatt, Titman, and Wermers 1997, Wermers 1999). Sias, Starks, and Titman (2006) develop a technique to extract the covariance of returns and institutional trading at a monthly or weekly level from quarterly data. They find that institutional trading leads returns. Griffin, Harris, and Topaloglu (2003), examine daily trading in Nasdaq 100 securities during a ten-month period beginning in May, 2000. They focus on the relation between returns and the buy-sell imbalance of individuals and institutions—not the total amount of trading within and between groups examined in this chapter.

3.1 Hypotheses

As discussed above, the low returns earned by stocks following high levels of buying by individuals could arise either from individuals pushing prices above fundamental value, or by institutions selling overvalued stocks to individuals. To differentiate between these alternatives, I develop and test a number of hypotheses.

While researchers typically think of liquidity provision as submitting a limit order that gives others the option to trade, Kaniel, Saar, and Titman (2008) note that practitioners think of a buy order placed when prices are falling, or a sell order placed when prices are rising, as supplying liquidity, regardless of whether the trader submits a limit or market order. This is the sense in which I use the term "liquidity provision" in the chapter. Individual investors may not set out to provide liquidity to institutions, actively posting limit buy and sell orders and taking the spread as compensation for their services; rather, they may respond to price changes caused by institutional trading and end up supplying liquidity. One way this can occur is if individuals have "latent" limit orders—prices at which they plan to buy or sell in the future—and these orders get triggered by price movements. For example, individuals who suffer from the disposition effect are more likely to sell a stock after seeing its price rise.[6] Grinblatt and Han (2005) show that momentum in stock returns can be caused by disposition-prone investors behaving this way.

[6]The disposition effect is the well-documented reluctance of investors to sell stocks below their purchase price. The behavior was first described by Shefrin and Statman (1985). See Seru, Shumway, and Stoffman (2007) and references therein.

Before stating the hypotheses, it is useful to consider possible price paths surrounding a trade, as shown in the stylized examples in Figure 3.1. The figure shows four price paths following a trade at time t_0. In the top two graphs, the trade is buyer-initiated. The bottom two graphs depict seller-initiated trades. The left two graphs show trade between an informed buyer and an uninformed seller, while the right two graphs show trade between an uninformed buyer and an informed seller. When the trade is initiated by an uninformed trader, prices subsequently revert, as seen in the northeast and southwest quadrants. If the trade initiator is informed, however, no such reversion takes place. This price reversion is a feature of models with asymmetric information: in contrast to the permanent price impact of informed trades, uninformed trading causes immediate price changes to compensate liquidity providers, but expected future cash flows have not changed.[7] The reversion can stem from bid-ask "bounce," and is critical to the estimation of liquidity measures such as Roll's (1984) spread and Pastor and Stambaugh's (2003) liquidity factor.

Given the poor performance of individual investors discussed above, the question studied in the chapter is whether individuals demand liquidity and actively move prices (top right of Figure 3.1) or supply liquidity as prices move (bottom left of the figure). If institutions are more likely to be informed, and individuals provide institutions liquidity in the sense defined above, then this stylized exam-

[7]See Glosten and Milgrom (1985), Kyle (1985), Easley and O'Hara (1987), Campbell, Grossman, and Wang (1993), Llorente, Michaely, Saar, and Wang (2002), Chordia and Subrahmanyam (2004), and Avramov, Chordia, and Goyal (2006), among others.

ple leads to a number of hypotheses. First, when institutions trade with individuals, prices should move. In particular,

Hypothesis 1. *When institutions purchase shares from individuals, prices contemporaneously increase. When institutions sell shares to individuals, prices contemporaneously decrease.*

Price increases accompanying institutional buying, and decreases accompanying institutional selling, are consistent with institutions demanding liquidity. In contrast, evidence against Hypothesis 1 would indicate that institutions supply liquidity. To test this hypothesis, I regress daily returns on a set of variables that summarize the amount of trading that took place between each investor group on each day. I also test the relation using weekly and monthly horizons. Details of the estimation procedure and results of tests of Hypothesis 1 are presented in Section 3.3.2.

Second, the stylized examples in the Figure 3.1 indicate that prices will change predictably after trading, depending on which types of investors caused the price change. In particular,

Hypothesis 2. *Price reversion is more likely following days when individuals trade with other individuals than days when individuals trade with institutions.*

Tests of this hypothesis are similar to those of Hypothesis 1, but instead of examining the contemporaneous relation between returns and trading by different groups, I investigate how returns change in the period following trading by individuals and institutions. In particular, I use a regression framework to test whether negative autocorrelation in daily returns is stronger following days when more trading takes place between two individuals than days when more trading occurs between individuals and institutions.

If Hypothesis 2 is true, it is also interesting to determine *whose* trading leads to price reversion. In particular, if trading comes primarily from individuals trading among themselves and prices change, we might expect institutions to react to the price movement by trading in a direction that pushes prices back to previous levels. That is, we would expect institutions to cause the price reversion by trading subsequently with individuals. This leads to the third hypothesis:

Hypothesis 3. *Institutions react to price changes caused when individuals trade with each other by subsequently trading with individuals to move prices back toward previous levels.*

To test Hypothesis 3, I examine the relation between institutional trading and the previous day's proportion of individual trading interacted with the price changes. I use a regression framework to test (a) whether institutions are more likely to *sell* to individuals following days that have both *high* returns and more intragroup individual trading; and (b) whether institutions are more likely to *buy* from individuals following days that have both *low* returns and more intragroup individual trading. Details of the estimation procedure and results for Hypotheses 2 and 3 are presented in Section 3.3.5.

3.2 Data and methods

In this section, I begin by describing the salient features of the data used in this study. I then discuss the procedures I use to classify investors into different groups, as this is key to the empirical implementation in the chapter.

3.2.1 Data description

The data used in this chapter are discussed in Chapter I. Summary statistics relevant for this chapter are presented in Table 3.2. The data include the transactions of nearly 1.3 million individuals and firms, beginning in January, 1995 and ending in December, 2003. For accounts that existed prior to 1995, opening account balances are also included, making it possible to reconstruct the total portfolio holdings of an account on each day. While the dataset includes exchange-traded options and certain irregular equity securities, I focus on trading in ordinary shares. After excluding trading in very thinly-traded securities,[8] there are more than 37.5 million trades in the data, including 10.7 million trades by 583,518 unique household accounts. The remaining trades are placed by financial institutions, corporations, and to a much smaller degree, government agencies and certain other organizations. Because all trading is recorded in the data, it is possible to construct measures of trader interaction that are not feasible with datasets that include small samples of the population.

[8] I exclude securities with fewer than 25,000 trades, or about ten trades per day during the sample period. This leaves 116 stocks. The additional requirement that daily returns data are available on Datastream reduces the sample to 106 stocks.

3.2.2 Complications

There are a number of complications that arise when working with these data especially relevant to this chapter. First, transactions are not time-stamped, so it is not possible to determine the order in which trades took place within a day. Second, while each transaction is identified as either a buy or a sell, the absence of quote data makes it impossible to identify which side initiated the trade, as is typically done with the quote-based algorithm of Lee and Ready (1991). Nevertheless, if prices consistently move in the same direction as an investor's trading, this investor must be initiating trades; an investor who sells shares into a rising market cannot be setting prices. To be more precise, since every transaction requires both a buyer and a seller, there is a sense in which both the buyer and the seller are setting prices. However, it is commonly understood in the literature that the initiator of a trade causes the price change; indeed, this is the identifying assumption of the Lee-Ready algorithm.

Third, the investor group classification is not always correct because of the existence of so-called "nominee" accounts. I outline my approach in dealing with these accounts in Section 3.2.2. A related issue is that financial institutions that act as market-makers are not separately identified in the data. In Section 3.2.2, I discuss the method I use to separate these accounts from my analysis.

Fourth, there is no direct match between the buy-side and sell-side of a transaction. For each trade, at least two observations are recorded in the data: one purchase and one sale. However, some trades are comprised of shares purchased or shares sold by more than one account, and in these situations there can be more than two records per trade. For example, if investor A buys 100 shares of Nokia and investor B sells 100 shares at the same price, they may have traded with each other, but no link between these transactions is reported in the data. This necessitates using a technique to identify the amount of trading that occurs between groups, which I discuss in Section 3.2.2.

Nominee accounts

Unlike investors domiciled in Finland, foreign investors are not required to register with the CSD. Instead, they may trade in an account that is registered in the name of a "nominee" financial institution. As a result, trading by foreign individuals can appear to be trading by the nominee institution. This is also true of American Depository Receipts (ADRs), which many Finnish firms have. For example, if an individual in the U.S. trades shares of Nokia's ADR on the NYSE, it will be classified in the CSD data as a trade by the institution that serves as a nominee for the ADR.

Since the focus of this chapter is on examining the differences in the price impact of trading among groups of investors, it is important to deal with this misclassification in some way. Note, however, that this misclassification makes it *more* difficult to find differences in the price impact of trading by different groups, because the misclassification causes a "mixing" of the group members. Nevertheless, I adopt the following approach to identify accounts that are likely to be nominee or ADR accounts. First, I require an institution to own at least ten percent of the total shares outstanding before considering it to be acting as a nominee. Then, for each stock and day, I calculate the number of trades and quantity of shares traded by each institution as a percent of all trading in that stock/day. I then count the number of days in a month in which the institution accounts for more than ten percent of both the number of trades placed and the quantity of shares traded. Finally, I classify as a nominee any institution with ten or more such days in a month. While the specifics of this procedure are somewhat arbitrary, it generally identifies only a few accounts that serve as nominees for each security, which is expected. Moreover, the results in this chapter are not sensitive to using a variety of alternative parameters of the classification procedure.

Market-makers

For the main analysis in this chapter, I focus on individuals and institutions that trade for information or liquidity reasons. Therefore, it is necessary to identify certain institutions and individuals that trade particularly actively, effectively acting as market-makers by both buying and selling a given stock on a particular day. While the motivation for such trading may differ, this type of trading may be broadly classified as market-making. (Linnainmaa (2007) explores the behavior of individual day traders in this market.) Traders who act as market-makers are unlikely to be trading based on fundamental or long-term information, and I therefore exclude this set of traders from my analysis.

To identify market-makers, I simply check whether an account purchased and sold shares in the same company on the same day. Occasionally, there are cases where an account buys and sells shares, but the amounts traded are different by orders of magnitude. For example, an account may purchase 5000 shares and sell 5 shares of one stock. It is unclear why this would occur, but it is probably not correct to call this trader a market-maker. Therefore, I require the amount purchased to be between 10 and 90 percent of the total amount traded by an account for it to be classified as a market-maker on that day. (The results are not sensitive to this restriction.) If an account is classified as a market-maker on at least five trading days in a particular month, I treat it as a market-maker for the entire month. The results in the chapter are unchanged if I exclude only accounts acting as market-makers on a day-by-day basis, or if I only allow financial institutions to be classified as market-makers. Note that market-makers are identified after the nominee accounts are identified, as discussed above in Section 3.2.2.

Identifying investor interaction

Given a classification of investors into groups, as in my data, it is possible to estimate the amount of trading that occurred between and within groups. For example, suppose trading on one day for one stock at one particular price is summarized as follows:

	Shares Bought	Shares Sold
Group A	250	450
Group B	2000	1800
Group C	250	250
Total	2500	2500

While we cannot be certain how much trade occurred between or within each group of investors, we may approximate these quantities by assuming that trade occurs in proportion to the amount of buying or selling accounted for by each group. Of the 250 shares purchased by Group A, we would therefore estimate that $450/2500 = 18\%$ (45 shares) were purchased from other members of Group A, $1800/2500 = 72\%$ (180 shares) were purchased from members of Group B, and so on. Continuing with this example, we would estimate the amount of trading within and between Groups A, B, and C as follows:

	Seller		
Buyer	A	B	C
A	45	180	25
B	360	1440	200
C	45	180	25

This estimation strategy conditions on the actual amount of purchases and sales by each group; we are *not* assuming that investors from each group randomly choose to buy and sell. Instead, we are assuming that, given the actual number of shares sold by a group, the probability of selling to any other group that bought shares is proportional to the relative amount of shares purchased and sold by each group.

This procedure is an approximation technique. In the previous example, it is possible that members of Group A did not purchase any shares from other members of Group A, although the estimate is that 45 shares were traded within this group. However, the procedure can yield an exact identification of the amount of trading that occurred between two groups. In the example, as is frequently the case in the data, one group (B) accounts for much of the purchases and sales. Since Group B accounts for so much of the trading, much of the trading must have occurred within this group—there simply is not enough selling by Groups A and C to meet the large demand for shares from Group B. In fact, by applying the procedure for *each price* at which the stock traded in a day and then aggregating to get a daily measure, I maximize the frequency with which this happens. Especially for all but the most frequently traded stocks, it is common for groups of investors to have only purchased or sold shares, but not both, at a particular price; it is therefore frequently possible to know with certainty exactly how much trading occurred between groups.[9] In my data, on average 28.2% (ranging from 3.4% to 60.3%) of a stock's trading volume is exactly identified in this manner.

An alternative method to determine how much trading occurs within and between groups involves establishing bounds on how much trading could have occurred. Referring again to the example, note that Group A investors could have traded no more than 250 shares with other Group A investors, since Group A only bought 250 shares in total. Similarly, intragroup trading in Group B could have been no more than 1800 shares. But since total purchasing by Groups A and C amounts to only 500 shares, Group B investors must also have sold no fewer than $1800 - 500 = 1300$ shares to other Group B investors. Therefore, intragroup trading for Group B must have been between 1300 and 1800 shares. Since any amount of trading within these bounds could be used as the estimated amount of trading, using the mean seems sensible. In the example, this value is quite close to the value estimated using the technique above—1550 vs. 1440, or 62% vs. 58% of total trading. When applied to my data, these two methods generally provide very

[9]Suppose that institutions purchased 800 shares and sold 1000 shares, and that individuals purchased 200 shares but sold no shares. The institutions must have sold 200 shares to the individuals and 800 shares to other institutions—there is no ambiguity here, and no estimation is required.

similar estimates of trade interaction, so the results reported below are generated from the first approach, which is somewhat less complicated. The results are not substantively changed, however, by using the second method.

3.3 Results

I turn now to presenting the results of the chatper. I begin by quantifying the amount of trading that occurs among investor groups. Institutions and individual investors account for a large proportion of total trading, and it is certainly possible that trading by either of these groups could have important prices effects. I then examine the relation between contemporaneous returns at a daily, weekly, or monthly horizon and inter- and intragroup trading quantities. Returns are higher when institutions buy shares from individuals, and lower when institutions sell shares to individuals, suggesting that institutions move prices. This indicates that, on average, institutions are informed and individuals are not. To allow returns and trading quantities to be mutually dependent, I also estimate a vector autoregression (VAR). The VAR allows me to estimate price impact functions for each of the groups in the data. Next, using transaction prices, I show that institutions buy stocks at higher prices than do individuals when returns are high, and sell at lower prices than do individuals when returns are low, which provides additional evidence at an intraday frequency that institutions set prices. I then show that price reversion is more prevalent when intragroup trading by individuals is high. Moreover, the price reversion is caused by the subsequent trading of institutions.

3.3.1 Investor interaction

Table 3.3 shows the amount of trading that is accounted for by each of the four largest groups of traders. The omitted groups are Government agencies, Nonprofit organizations, and Registered foreigners; for most stocks on most days, these groups account for a negligible amount of total trading. The total amount of trading may be calculated as a percent of the number of shares traded (presented in Row 1) or as a percent of the value of shares traded (Row 2). The combined trading of individuals, institutions, nominees, and other corporations accounts for approximately 95% of all trading. About 40% comes from institutions, and about 20% from individuals ("Households"). Individuals account for 20.6% of trading by number of shares traded, and 17.4% by value, indicating that individuals tend to trade more heavily in low-priced stocks.

Using the total proportion of trade from each group, it is possible to calculate the amount of trading we would expect to occur within and between each group. The results of this exercise, using the percent of trading by shares (Row 1) is reported in the table under the heading "Expected % of all trading." For example, since financial institutions account for 38.1% of all trading, $0.381^2 = 14.5\%$ of all trading should be between two institutions. Similarly, the amount of trading that would occur between institutions and individuals is $0.381 \times 0.206 \times 2 = 0.157$, where we multiply by two because either the institution can buy from the individual or the individual can buy from the institution.[10]

The lower panel of Table 3.3 shows the estimated amount of trading that occurred between and within groups, using the technique described above in Section 3.2.2. That is, in contrast to the "expected" percentage in the top panel, we now condition on the actual amount of buying and selling by each group on each day, rather than total trading volume alone. There are several notable differences between actual and expected trading. Institutions trade more with each other, and less with individuals, than expected. As well, individuals trade more with other individuals than expected. Previous research has documented herding behavior among both individuals and institutions; in other words, within-group trading is positively correlated. This means that trading between groups should be quite common—if all individuals are buying, they cannot trade with each other. The results presented here show that the previous findings mask an important fact: a great deal of trading occurs between two individuals or between two institutions. In these data, individuals are more than twice as likely to trade with other individuals than their trading volume would suggest, and only 9.3% of trading takes place between individuals and institutions. Alternatively, ignoring nonfinancial corporations and ADRs, we see that institutions are about half as likely to trade with individuals as they are to trade with other each other.

[10]The table is lower-diagonal because it shows the expected amount of total trading—not separated into buying and selling. A matrix showing the expected proportion of buying and selling separately would be symmetric, with the off-diagonal elements equal to one-half of the reported off-diagonal elements.

What explains the discrepancy between the actual and expected amount of trading among individuals? There are several possibilities, each of which could partially contribute to the result. First, institutions can arrange large block trades with each other away from the regular limit-order book, so for a fixed amount of trading, less volume will take place between individuals. Second, as in the model of Easley and O'Hara (1987), informed investors might only trade when there has been an information event. If institutions tend to be informed, they will be less likely to trade when no information event has occurred, and any trading by individuals will tend to be with other individuals. Indeed, Easley, Engle, O'Hara, and Wu (2002) find that uninformed orders are clustered in time, but also that uninformed investors avoid trading when informed investors are likely to be present. A related explanation is offered by Barber and Odean (2006), who document that individual investors are more likely to trade following events that get media attention.

It is worth noting at this point that the amount of trading coming from each group determines how much power I will have to find a relation between group trading and price changes. For example, if almost all trading came from institutions, then it would be difficult to find a relation between price changes and the trading of any group because returns (the left hand side variable in my regressions) would vary, but the proportion of trading from each group (the right hand side variable) would not. The percent of trading reported in Table 3.3 suggests that trading is sufficiently spread among different groups to provide adequate power for my tests.

The results in Table 3.3 are calculated using data aggregated from all stocks in the sample. In the Finnish market, a few stocks account for much of the trading volume and market capitalization. Notably, the largest stock, Nokia, makes up 36% of the total stock market capitalization on average during the sample period (ranging daily from 16% to a high of 64% at one point in 2000). On average, Nokia accounts for 57% of the value of daily trading (ranging from 3% to 95%). To confirm that the reported results apply to Finnish stocks in general and are not driven by Nokia or a few other large stocks, I first calculate a time-series average of the daily interaction estimates for each stock, and then examine cross-sectional statistics for these means. This procedure puts equal weight on each stock's average interaction estimates. These results are reported in Table 3.4. Across the 106 stocks in my sample, trading among institutions and among individuals accounts for an average of 14.7% and 19.3%, respectively, confirming that intragroup trading is important across stocks. Trading among individuals ranges from 0.1% to 60.6% of all trading; the smaller numbers come from large actively-traded stocks in which trading is dominated by institutions, while the larger numbers come from smaller, less-liquid stocks. It is interesting to note that the range for trading between institutions and individuals is much narrower, indicating that it is generally an important part of trade volume, although it never accounts for more than 14.8% of the total.

The data presented in this section show that individuals account for approximately one-fifth of trading in the Finnish stock market. This is certainly sufficiently large for the trading of individual investors to have substantial price effects. About half of individual investors' trading is with other individuals, and half is with institutions. In the next section, I examine which type of trading is most associated with price changes.

3.3.2 Daily returns

To understand how prices are determined by the interaction of different investors in the market, I examine the relation between returns and the proportion of trading within and between investor groups. In particular, I estimate the regression

$$\begin{aligned} R_{i,t} = \alpha + \beta R_{i,t-1} + \gamma_1 \text{Inst/Inst}_{i,t} + \gamma_2 \text{Inst/Ind}_{i,t} \\ + \gamma_3 \text{Ind/Inst}_{i,t} + \gamma_4 \text{Ind/Ind}_{i,t} + \epsilon_{i,t}, \quad (3.1) \end{aligned}$$

where the notation $\text{A/B}_{i,t}$ represents the proportion of trading that is accounted for by investors from Group A purchasing shares from investors in Group B. "Ind" and "Inst" denote individuals and institutions, respectively. For brevity, I focus in the remainder of the chapter only on the trading of individuals and institutions.[11] Since not all investor groups are included in the regression, the A/B trade variables do not sum to one, and including an intercept does not result in perfect collinearity.

[11]The main conclusions are not altered if I include the other investor groups. In particular, the price impact estimates for nonfinancial corporations and nominee accounts lie between those of individuals and institutions, and the estimated coefficients for individuals and institutions are basically unchanged if these additional groups are included.

I begin by estimating the regression at a daily horizon, and later confirm that similar results obtain using weekly or monthly data. If trading between members of Group A and Group B (possibly the same group) leads to price changes, then contemporaneous returns will be positive or negative, and the estimated γ coefficient for the relevant combination of trade will be significant. If no group has a consistent effect on prices, then the four trade variables will be economically small and statistically insignificant. The lag return, $R_{i,t-1}$, is included to control for the known autocorrelation in daily returns.

I estimate regression (3.1) using the approach of Fama and MacBeth (1973) in two ways: averaging results from time-series regressions by stock or from cross-sectional regressions by date. The latter approach results in a time-series of estimates, which may be autocorrelated. Therefore, I use the robust standard error calculation of Newey and West (1987) with five lags, which corresponds to one week of trading.[12] The adjustment for autocorrelation is unnecessary for the cross-section of estimates obtained from the first approach. Estimating the regression separately along each dimension allows me to check whether the results are driven by a cross-sectional relation, a time-series relation, or both. Using a panel regression with clustered standard errors yields similar results, but could suffer from a bias due to the lagged dependent variable, so I do not report these estimates.

[12] With 2184 daily observations, the rule of thumb noted by Greene (2005, p. 267) calls for $\lfloor \sqrt[4]{2184} \rfloor = 7$ lags, and the significance level of all results is unchanged by using the longer lag length.

Coefficient estimates for regression (3.1) are presented in Table 3.5, and provide an interesting picture of how prices are affected by the interaction of individual and institutional investors in the market. Panel A reports results where the A/B trade variables are calculated as the proportion of total trade volume each day, and Panel B uses trade variables calculated as the proportion of daily turnover (shares traded divided by number of shares outstanding) each day. Results from the stock-by-stock regressions are denoted "FM by stock," and the day-by-day regressions are denoted "FM by date."

Consider first the "FM by stock" regressions in Panel A. When institutions trade with other institutions, there is no significant change in price. As discussed above in Section 3.2.2, a large portion of trading comes from institutions trading with other institutions and households trading with other households. To date, the literature has not been able to address the importance of price changes that occur during this trading. The results in Table 3.5 bear directly on this question. The estimated coefficient is 0.0092, with a t-statistic of 1.35. Of the 106 stock-by-stock regressions, the Inst/Inst coefficient is significantly negative in 8% of regressions and significantly positive in 11% of regressions, using a 5% significance level. Similarly, trading between individuals does not yield significant results, although 27% of stocks yield statistically negative coefficients. That is, there appears to be little in the way of price changes when trading takes place within each investor group.

In sharp contrast, when institutions purchase from individuals, prices increase. The estimated coefficient of 0.0316 is highly significant, and fully 61% of the regressions yield statistically positive results; none of the coefficients is statistically negative. And when individuals buy from institutions, prices fall: the coefficient of -0.0423 is again highly significant, and 72% of the cross-sectional regressions yield statistically positive coefficients, with none significantly negative. The cross-sectional average of the stock-level standard deviation in the trade variables is reported in the table under the heading "C.S. Std. Dev." These values can be used to ascertain the economic significance of the coefficient estimates. A one standard deviation increase in the Inst/Ind variable (that is, institutions buying shares from individuals) increases daily returns by $0.316 \times 0.1151 = 0.0036$ or 36 basis points (bps). This is economically large relative to an average daily return of approximately 10 bps. Similarly, a one standard deviation increase in the Ind/Inst variable leads to a decrease in returns of $0.0423 \times 0.1145 = 48$ bps.

The day-by-day cross-sectional ("FM by date") regressions deliver similar results, suggesting that the result is not driven solely by either time-series or cross-sectional relations. The effect of trading between households in this specification is to reduce contemporaneous returns. This result is not robust, however, which can be seen by looking at the results in Panel B. Across both panels and both specifications, the result that is consistent is that prices increase when individuals sell shares to institutions, and prices decrease when individuals buy shares from individuals. This evidence strongly supports Hypothesis 1.

3.3.3 Vector autoregressions

Another approach to examining the relation between group trading and subsequent price changes is to estimate a vector autoregression (VAR) as in Hasbrouck

(1991). There are a number of benefits to this approach. First, allowing the trade variables and returns to depend on lags of each other provides a way to examine potentially complicated dynamics among the variables. Second, the lag structure of the VAR allows me to plot contemporaneous price impact and subsequent price changes. These plots are the empirical analogue of the stylized price paths shown in Figure 3.1.

Let $\mathbf{y}_{i,t} \equiv (\text{Ind/Ind}_{i,t}, \text{Inst/Inst}_{i,t}, \text{Ind/Inst}_{i,t}, \text{Inst/Ind}_{i,t}, R_{i,t})'$. As above, the notation $\text{A/B}_{i,t}$ denotes the proportion of trading in stock i on date t that comes from Group A purchasing shares from Group B, and "Ind" and "Inst" denote individuals and institutions, respectively. The reduced-form VAR for each stock, i, is

$$\mathbf{y}_{i,t} = \sum_{k=1}^{p} \mathbf{\Phi}_k \mathbf{y}_{i,t-k} + \boldsymbol{\epsilon}_{i,t}. \tag{3.2}$$

In order to allow returns to depend contemporaneously on the trade variables, I estimate a dynamic structural VAR (see Hamilton (1994, Section 11.6)). In particular, triangular factorization of the error covariance matrix, $\mathbf{\Sigma} \equiv E(\boldsymbol{\epsilon}_t \boldsymbol{\epsilon}_t')$, yields a lower-diagonal matrix, $\mathbf{A}_0$, with ones on the principal diagonal such that $\mathbf{A}_0 \mathbf{\Sigma} \mathbf{A}_0' = \mathbf{\Sigma}^d$ where $\mathbf{\Sigma}^d$ is a diagonal matrix with all positive elements. Multiplying both sides of equation (3.2) by $\mathbf{A}_0$ gives the dynamic structural VAR

$$\mathbf{A}_0 \mathbf{y}_{i,t} = \sum_{i=1}^{p} \mathbf{A}_k \mathbf{y}_{i,t-k} + \boldsymbol{\eta}_{i,t}, \tag{3.3}$$

where $\mathbf{A}_k = \mathbf{A}_0 \mathbf{\Phi}_k$ and $\boldsymbol{\eta}_t = \mathbf{A}_0 \boldsymbol{\epsilon}_t$. The shocks in this system are uncorrelated, since $E(\boldsymbol{\eta}_t \boldsymbol{\eta}_t') = E(\mathbf{A}_0 \boldsymbol{\epsilon}_t \boldsymbol{\epsilon}_t' \mathbf{A}_0') = \mathbf{\Sigma}^d$. Moreover, since $\mathbf{A}_0$ is lower-diagonal, this specification allows each variable in $\mathbf{y}_{i,t}$ to depend on contemporaneous realizations of the variables that precede it in the vector:

$$\text{Ind/Ind}_{i,t} = \sum_{i=1}^{p} \mathbf{A}_{1k} \mathbf{y}_{t-1} + \eta_{i,1t}, \tag{3.4a}$$

$$\text{Inst/Inst}_{i,t} = \sum_{i=1}^{p} \mathbf{A}_{2k} \mathbf{y}_{t-1} - a_{21} \text{Ind/Ind}_{i,t} + \eta_{i,2t}, \tag{3.4b}$$

$$\text{Ind/Inst}_{i,t} = \sum_{i=1}^{p} \mathbf{A}_{3k} \mathbf{y}_{t-1} - a_{31} \text{Ind/Ind}_{i,t} - a_{32} \text{Inst/Inst}_{i,t} + \eta_{i,3t}, \tag{3.4c}$$

$$\text{Inst/Ind}_{i,t} = \sum_{i=1}^{p} \mathbf{A}_{4k} \mathbf{y}_{t-1} - a_{41} \text{Ind/Ind}_{i,t} - a_{42} \text{Inst/Inst}_{i,t} - a_{43} \text{Ind/Inst}_{i,t} + \eta_{i,4t}, \tag{3.4d}$$

and

$$\begin{aligned} R_{i,t} = \sum_{i=1}^{p} \mathbf{A}_{5k} \mathbf{y}_{t-1} &- a_{51} \text{Ind/Ind}_{i,t} - a_{52} \text{Inst/Inst}_{i,t} \\ &- a_{53} \text{Ind/Inst}_{i,t} - a_{54} \text{Inst/Ind}_{i,t} + \eta_{i,5t}. \end{aligned} \tag{3.4e}$$

Here a_{mn} denotes the (m, n)th element of $\mathbf{A}_0$, and $\mathbf{A}_{jk}$ denotes the jth row of $\mathbf{A}_k$. Since the order in which the variables appear in the $\mathbf{y}_{i,t}$ vector determines which variables are allowed to affect other variables contemporaneously, there is a strong theoretical reason to put returns last, but the order of the trade variables is less obvious. I therefore confirm that all the results below are unaffected by permuting the order of the trade variables.

Methods to estimate VARs in panel data are not well-developed. Therefore, in the spirit of Fama and MacBeth (1973), I separately estimate the dynamic structural VAR for each stock and then take cross-sectional means of coefficient estimates. Statistical significance is determined from the cross-sectional standard errors of these means. I choose ten lags ($p = 10$) by examining the Akaike Information Criterion for the VAR. While a lower-order VAR fits well for some stocks, I fit the same model to all stocks to ease comparison of results. Estimating a model with five lags yields results that are substantially the same as those reported here. Estimation of (3.2) yields estimates of the $\mathbf{\Phi}_k$ matrices, and triangular factorization of the estimated error covariance matrix gives an estimate of $\mathbf{A}_0$, which is then used to calculate estimates of the $\mathbf{A}_k$ coefficient matrices. This is repeated for each of the 106 stocks in the sample, and Table 3.6 summarizes the results from these regressions for $k = 0, 1, 2$. For brevity, higher-order lags and the constant term are not reported.

Controlling for complex serial correlations does not alter the results reported above. The contemporaneous effect ($k = 0$) on returns of the Ind/Inst and Inst/Ind variables are very close to those reported in the previous table. In this specification, there is also evidence that returns are negative when individuals trade with each other: the Ind/Ind variable is negative and significant. The magnitude of this effect is considerably smaller than the effect of intergroup trading.

Looking at $k = 1$, the negative autocorrelation in daily returns that is consistent with bid-ask bounce is quite prevalent in these data. The amount of within- and between-group trading is positively autocorrelated up to three lags. For example, at the first lag ($k = 1$), the autocorrelation coefficients range from 0.0964 for the Inst/Inst variable to 0.1493 for the Ind/Inst variable, and all are significant at the 1% level. This pattern holds for higher-order lags and for all of the variables except returns.

Empirical price paths

The coefficient estimates from the VAR can be used to construct impulse response functions, which are the empirical analogue to the stylized price paths presented earlier in Figure 3.1. That is, I calculate the effect of a one standard deviation impulse to each of the elements of the orthogonalized shocks, $\boldsymbol{\eta}_{i,t}$, one at a time. For example, to see the effect on returns of a shock to Inst/Ind I use equations (3.4d) and (3.4e) to estimate the increase in returns caused by a one standard deviation increase in $\eta_{i,4t}$.

Results from applying this procedure to the return equation of the VAR are presented in Figure 3.2. The figure plots the price impact function for the variables Inst/Ind (solid line), Inst/Inst (short dash), Ind/Ind (long dash), and Ind/Inst (dash-dot). Time is measured in trading days, so the ten lags that are plotted correspond to two weeks of trading. The graph begins at time -1 to show the price impact in the first day relative to the previous day's close. A one standard deviation innovation in Inst/Ind increases the contemporaneous return by 21 bps, and there is no evidence of subsequent reversion. Similarly, a one standard deviation innovation in Ind/Inst leads to a return of -28 bps, and there is no subsequent reversion. In sharp contrast, a one standard deviation innovation in Ind/Ind is associated with a return of -12 bps, but prices subsequently revert. Comparing these empirical price impact functions to the stylized examples in Figure 3.1 indicates that when institutions purchase shares from individuals or sell shares to individuals they are informed traders demanding liquidity from individuals. This result is clearly not consistent with individual investors actively moving prices.

3.3.4 Alternate horizons

Weekly and monthly results

Table 3.7 presents results for estimation of regression (3.1) over different trading horizons. Panel A shows results when the trading percentage variable and returns are calculated over weekly horizons, and Panel B shows results calculated at a monthly horizon. At these longer horizons, the results remain consistent with what was found in the daily regression. Daily returns are contemporaneously higher when institutions purchase shares from individuals and lower when they sell shares to individuals. There is little or no price effect from intragroup trading by individuals or institutions. As in the daily results, the price impact of institutions is stronger when they sell to individuals than when they buy from individuals.

It is interesting to note that the magnitude of the effect on returns (in absolute terms) is greater when institutions sell to households than when institutions buy from households, especially at the monthly horizon. That is, institutions appear to have larger price impact when they sell to individuals than when they buy from individuals. Campbell, Ramadorai, and Schwartz (2008) find a similar asymmetry, and suggest that it could stem from the reluctance of institutions to use short sales.[13] This asymmetry is not apparent in the daily results presented earlier, perhaps because daily returns are small and can be affected by institutions with small positions, but short-sale constraints then become binding at longer horizons.

Intraday results: Evidence from trading prices

A potential concern with the daily results presented above relates to the timing of trades within the day. It is possible that trading by individuals moves prices, and that institutions subsequently trade at those new prices, but that institutional trading does not actually affect prices. For example, suppose individuals trade in the morning at prices above the previous day's close, and when institutions see prices increasing they act as momentum traders and decide to buy shares. Their buying, however, could occur at prices that are not higher than the prices set by individual trading. In this situation, we would find that institutions purchase shares on days when prices rise, but institutions did not cause the price change.

The strength and robustness of the results above suggest that this scenario is unlikely. When individuals purchase shares, same-day returns are typically lower than when they sell shares. Nevertheless, an additional test to rule out this possibility is in order. Unfortunately, the transactions in the dataset are not time-

[13]Their argument is as follows: if institutions wish to increase their exposure to a particular risk factor, they may buy stocks that load strongly on that factor. If their buying causes price increases in one stock, they can purchase shares of another stock that also loads on the factor, thereby reducing their overall price impact. However, if an institution wishes to decrease its exposure to a factor and is reluctant or unable to use short sales, it can only sell stocks it currently owns. This forces the institution to sell more aggressively the stocks it owns, and could lead to larger price impact.

stamped, so it is not possible to examine the order in which trades were placed and the path prices took within the day. However, we do observe trade prices for each transaction, so it is possible to compare the prices at which institutions and individuals purchase and sell shares, and the relation between these prices and contemporaneous returns.

To understand this test, suppose that on a particular day, a stock trades only at two prices, \$10 and \$11. Suppose further that both Groups A and B bought at \$10 and sold at \$11, but only A bought at \$11. If the closing price is \$11, it can only be because Group A moved the price; Group B did not purchase any shares at \$11, so they could not have caused the price to move up to that level. That is, since Group B bought at a lower price than did Group A—and prices increased—it is not possible for Group B to have caused prices to move. This suggests that we can test whether one group moves prices by examining the relation between returns and the difference in the purchase or sale prices of individuals and institutions. The point of this exercise is *not* to determine trading profits within a day, since we are not comparing one group's purchase and sale prices; rather, we are looking at whether one group purchased stocks at higher prices than did another group on days when prices rose, or sold stocks at lower prices on days when prices fell.

This test also allows us to differentiate intraday price movements from close-to-open price movements. For example, suppose the opening price is above the previous day's close, but then remains flat during the day. If most trading just happened to come from institutions buying shares from individuals, the earlier regressions would find that institutions buy from individuals when prices increase—

but the price change happened entirely when the market was closed. This intraday test, however, would find that individuals and institutions both bought at the same price on a day when returns increased.

Of course, stocks typically trade at more than two prices on one day, so I look instead at the average price at which institutions and individuals buy or sell shares. Specifically, I test the relation using Fama-MacBeth estimation of the regressions

$$\bar{b}^I_{i,t} - \bar{b}^H_{i,t} = \beta_1 \cdot \mathbf{1}_{\{R_{i,t}>0\}} + \beta_2 \cdot \mathbf{1}_{\{R_{i,t}<0\}} + \epsilon_{i,t} \tag{3.5}$$

and

$$\bar{s}^I_{i,t} - \bar{s}^H_{i,t} = \beta_1 \cdot \mathbf{1}_{\{R_{i,t}>0\}} + \beta_2 \cdot \mathbf{1}_{\{R_{i,t}<0\}} + \epsilon_{i,t}, \tag{3.6}$$

where $\bar{b}$ and $\bar{s}$ denote the average purchase and sale prices, respectively, for stock i on date t, and superscript I and H denote institutions and households, respectively. The indicator function, $\mathbf{1}_{\{\cdot\}}$, takes a value of one only when the condition in curly brackets is true; otherwise it is zero. Intercepts are excluded to prevent perfect collinearity. In a second set of regressions, I allow the purchase or sale price difference to vary with the magnitude of the return:

$$\bar{b}^I_{i,t} - \bar{b}^H_{i,t} = \alpha + \beta_1 \cdot \mathbf{1}_{\{R_{i,t}>0\}}|R_{i,t}| + \beta_2 \cdot \mathbf{1}_{\{R_{i,t}<0\}}|R_{i,t}| + \epsilon_{i,t} \tag{3.7}$$

and

$$\bar{s}^I_{i,t} - \bar{s}^H_{i,t} = \alpha + \beta_1 \cdot \mathbf{1}_{\{R_{i,t}>0\}}|R_{i,t}| + \beta_2 \cdot \mathbf{1}_{\{R_{i,t}<0\}}|R_{i,t}| + \epsilon_{i,t}. \tag{3.8}$$

Excluding intercepts in this second set of regressions is unnecessary. In each of these regressions, it is important to scale the prices so that the magnitude of the estimates may be compared across stocks, which is key to the Fama-MacBeth approach. Therefore, I scale the price difference by each stock's average trade price for that day. Scaling by the average closing price in the previous week yields similar results, as does using share-weighted prices instead of a simple average.

Estimates of regressions (3.5)–(3.8) are presented in Table 3.8. The results are not consistent with individuals moving prices. Focusing on Column 1, on days when returns are positive, institutions purchase shares at higher prices than individuals purchase shares, but on days when prices are negative, there is no statistical difference between the purchase prices. Economically, there is clearly a large difference between the coefficients on positive- and negative-return days (25 bps compared to 2 bps). Similarly, on days when returns are negative, institutions sell at lower prices than do individuals (Column 3). On the sell side, the economic magnitude of the difference is not as large as it is for purchase prices, but there is still a very large statistical difference. Columns 2 and 4 present results for similar regressions, but here we allow the effect size to vary with the level of returns. The statistical differences are even stronger in this specification. The economic magnitude of these estimates is not particularly large: when a stock has a positive return of 1%, institutions purchase shares at about 6.5 bps above individual purchase prices. Nevertheless, the magnitude of the difference between positive and negative return coefficients is both economically and statistically large.

These results convincingly show that prices move in response to institutional demand. Price changes cannot be attributed to the trading of individuals: when prices rise, institutions buy shares at higher prices than do individuals, and when prices fall, institutions sell shares at lower prices than do individuals. In either case, prices are clearly not being pushed up or down by individuals. Combined with the evidence presented in Section 3.3.2 for daily, weekly, and monthly horizons, these results provide strong support for Hypothesis 1.

3.3.5 Returns following trade

The second hypothesis to be tested is that price reversion is more likely following days when trading is dominated by intragroup individual trading. It is possible that when the bulk of trading is between individuals, without much institutional trading, prices are pushed away from fundamental values. If this is the case, we might expect prices to revert in subsequent trading. To examine this, I estimate the regression

$$\begin{aligned} R_{i,t} = \alpha &+ \beta_1 \text{Ind/Ind}_{i,t-1} R_{i,t-1} + \beta_2 \text{Inst/Inst}_{i,t-1} R_{i,t-1} \\ &+ \gamma_1 R_{i,t-1} + \gamma_2 \text{Ind/Ind}_{i,t-1} + \gamma_3 \text{Inst/Inst}_{i,t-1} \\ &+ \gamma_4 \text{Ind/Inst}_{i,t-1} + \gamma_5 \text{Inst/Ind}_{i,t-1} + \epsilon_{i,t}, \end{aligned} \tag{3.9}$$

for stock i on date t. “Inst” denotes financial institutions, and “Ind” denotes individual investors. Pairs of groups represent purchasing of shares by the first group from the second group. For example, "Ind/Inst" is the proportion of trading accounted for by individuals buying shares from institutions.[14]

If returns tend to revert after days when individuals trade with other individuals and returns are either high or low, β_1 should be negative. As shown in the results presented in Table 3.9, this is precisely what we find. The negative relation appears in both the cross-section and time-series Fama-MacBeth regressions. Moreover, there is no such effect for intragroup institutions trading—β_2 is not significantly different from zero in either set of regressions. Consistent with the VAR results presented above, the estimates in the last two lines of Table 3.9 show that days when individuals buy shares from institutions—which we previously found to be days when prices decline—are followed by further price declines. And when institutions buy shares from individuals—which are days when prices increase—prices tend to increase further on the following day. That is, prices revert less than usual when there is more intragroup trading; they revert more when individuals trade with other individuals or when institutions trade with other institutions. These results strongly support Hypothesis 2, that price reversion is more likely after intragroup trading by individuals than after trading between the groups.

The price reversion that we observe must be caused by trading between or within the groups. In particular, it is possible that institutions react to the price movements caused by individuals by subsequently purchasing (selling) underpriced (overpriced) shares from individuals. To investigate this, I examine the

[14]In the remainder of the chapter I use level of the trade variables and not the unexpected components from the VAR in section 3.3.3. All reported results remain unchanged if I instead use the residuals from each of the trade variable equations.

relation between institutional trading with individuals on date t and intragroup trading by individuals on date $t-1$. Specifically, I estimate the regressions

$$\begin{aligned}\text{Inst/Ind}_{i,t} = \alpha &+ \beta_1 \text{Ind/Ind}_{i,t-1} R_{i,t-1} + \beta_2 \text{Inst/Inst}_{i,t-1} R_{i,t-1} \\ &+ \gamma_1 R_{i,t-1} + \gamma_2 \text{Ind/Ind}_{i,t-1} + \gamma_3 \text{Inst/Inst}_{i,t-1} \\ &+ \gamma_4 \text{Ind/Inst}_{i,t-1} + \gamma_5 \text{Inst/Ind}_{i,t-1} + \epsilon_{i,t}, \quad (3.10)\end{aligned}$$

and

$$\begin{aligned}\text{Ind/Inst}_{i,t} = \alpha &+ \beta_1 \text{Ind/Ind}_{i,t-1} R_{i,t-1} + \beta_2 \text{Inst/Inst}_{i,t-1} R_{i,t-1} \\ &+ \gamma_1 R_{i,t-1} + \gamma_2 \text{Ind/Ind}_{i,t-1} + \gamma_3 \text{Inst/Inst}_{i,t-1} \\ &+ \gamma_4 \text{Ind/Inst}_{i,t-1} + \gamma_5 \text{Inst/Ind}_{i,t-1} + \epsilon_{i,t}, \quad (3.11)\end{aligned}$$

for stock i on date t. The dependent variable in regressions (3.10) and (3.11) is either $\text{Inst/Ind}_{i,t}$ or $\text{Ind/Inst}_{i,t}$—that is, the proportion of trading accounted for by institutions purchasing shares from individuals or the proportion accounted for by individuals purchasing shares from institutions.

Suppose trading on date $t-1$ came largely from individuals trading with other individuals, and that returns were positive. If this trading moved prices above fundamentals, then we would expect institutions to be less likely to buy shares from individuals on date t and more likely to sell shares to individuals on date t. That is, we would expect β_1 to be negative in regression (3.10) and positive in regression (3.11). As shown in the estimation results presented in Table 3.10, this prediction is borne out by the data. Institutions are less likely to buy shares from individuals (Panel A) and more likely to sell shares to individuals (Panel B) if more trading on $t-1$ occurred between two individuals and prices increased. Combined with the results in Table 3.9, these results indicate that if prices are moved by individuals, institutions subsequently trade with individuals in a direction that leads prices to revert. This evidence provides strong support for Hypothesis 3.

3.4 Conclusion

This chapter studies the relation between returns and trading by individual and institutional investors. I show that trading between individuals and institutions is relatively uncommon, but that these trades consistently lead to price changes. In contrast to the recent work of Barber, Odean, and Zhu (2006) and Hvidkjaer (2006), I find very little evidence of price effects from the trading of individual investors. In addition, I show that when prices change as a result of individual investors trading with other individuals, price reversion is more common than when they trade with institutions. Moreover, this reversion is caused by institutions subsequently trading with individuals in a direction that pushes prices back toward previous levels.

There are a number of reasons why it is important to understand what type of trading leads to price changes. First, if prices are determined primarily by traders with poor information or are contaminated by beliefs of investors suffering from behavioral biases, then features of the return series, such as volatility and covariance with macroeconomic variables like consumption may be spurious. This is of obvious importance to research in asset pricing, which seeks to explain relations among these variables. Second, a large literature seeks to understand the performance of mutual funds. Part of the challenge in assessing this performance is the difficulty of determining whether trading by mutual funds causes price movements. I find that the trading of mutual funds and other institutions moves prices, so findings of positive correlation between institutional trading and returns at low frequencies cannot be taken as evidence of good performance. However, my results do not rule out the possibility of further price movements allowing funds to earn profits.

In contrast to previous papers, the results presented here are calculated with daily data. At this relatively high frequency, it is apparent that price movements are generally caused by the trading demands of institutions. And if prices do move when individuals trade with each other, they quickly revert. This suggests that while there may be short-term price effects caused by individual investors, prices are unlikely to be affected by such distortions at longer horizons.

Table 3.1: Stock Returns Following Trading by Institutions and Households

This table presents average returns to stocks in the year subsequent to trading by institutions and individuals, separated by the rank of each group's buying pressure. In June and December of each year, firms are assigned to one of three portfolios for each group based on the proportion of buy volume for that group, $B_{i,t}/(B_{i,t}+S_{i,t})$. The table reports the returns to these portfolios over the subsequent year. t-statistics are presented in parentheses. Statistical significance at the 10%, 5% and 1% levels is denoted by †, *, and **, respectively.

	Institutions	Households
Most sold	-0.0011	0.1175
	(-0.04)	(4.34)**
Most bought	0.0892	0.0492
	(3.12)**	(1.85)†

Table 3.2: Summary Statistics

This table provides summary statistics for the data. The sample period is 1995–2003. The "Nominees/ADRs" group is identified using the technique explained in Section 3.2.2. "Other" includes Government Agencies, Nonprofit Organizations, and Registered Foreigners. The number of accounts and total number of trades are shown in the first two columns, respectively. The remaining columns present averages for per-trade, per-day, or per-account data.

	Number of:		Average of:			
	Accounts	Trades (MM)	Shares per trade (000s)	Value per trade (000s)	Trades per day (000s)	Securities traded
Financial institutions	920	14.0	3.6	77.8	6.4	33
Households	583,518	10.7	1.2	7.3	4.9	4
Nominees / ADRs	47	7.4	4.2	65.5	3.4	169
Nonfinancial corporations	29,186	4.6	3.8	57.3	2.1	8
Other	12,269	0.8	1.6	30.8	0.4	2
All	625,940	37.5	2.9	47.8	3.4	43

Table 3.3: Trader Interaction

This table shows the amount of trading (in percent) that occurs in total and between the different investor groups. The expected amount of trading between groups is calculated assuming random interaction in proportion to the percent of total trades based on number of shares traded presented in the first row. The actual amount of trading is estimated using the technique discussed in the text. The total number of shares traded between each group is calculated for each stock/day during the period 1995–2003. The percentages do not sum to 100% because the trading of certain groups that account for very little total trading volume is omitted.

	Financial Institutions	Households	Nominees / ADRs	Nonfinancial Corporations
Percent of trading (by shares)	38.1	20.6	15.2	20.9
Percent of trading (by value)	40.9	17.4	16.1	20.5
Expected % of trading with:				
Financial institutions	14.5			
Households	15.7	4.2		
Nominees / ADRs	11.6	6.2	2.3	
Nonfinancial corporations	15.9	8.6	6.4	4.4
Estimated % of trading with:				
Financial institutions	20.6			
Households	9.3	9.1		
Nominees / ADRs	10.1	3.5	5.7	
Nonfinancial corporations	13.9	8.1	4.9	6.2

Table 3.4: Trader Interaction—Cross-Sectional Statistics

This table shows the average amount of trading (buying and selling, in percent) that occurs between the different investor groups. The amount of trading is estimated for each stock/day using the technique discussed in the text. The time-series average is then calculated for each stock, and the table reports cross-sectional statistics for those means ($N = 106$ stocks).

Group 1	Group 2	Mean	Std. Dev.	Min.	Max.
Financial institutions	Financial institutions	14.7	10.5	0.5	41.4
	Households	5.9	2.6	1.6	14.8
	Nominees / ADRs	4.0	3.8	0.0	14.4
	Nonfinancial corporations	5.2	2.6	0.3	13.2
Households	Households	19.3	17.1	0.1	60.6
	Nominees / ADRs	2.6	1.9	0.0	8.6
	Nonfinancial corporations	6.1	4.4	0.5	17.1
Nominees / ADRs	Nominees / ADRs	4.2	4.5	0.0	26.4
	Nonfinancial corporations	1.6	1.1	0.0	4.6
Nonfinancial corporations	Nonfinancial corporations	4.5	2.2	0.5	15.9

Table 3.5: Returns and Group Interaction—Daily

This table presents the results of the regression

$$R_{i,t} = \alpha + \beta R_{i,t-1} + \gamma_1 \text{Inst/Inst}_{i,t} + \gamma_2 \text{Inst/Ind}_{i,t} + \gamma_3 \text{Ind/Inst}_{i,t} + \gamma_4 \text{Ind/Ind}_{i,t} + \epsilon_{i,t},$$

where the notation $\text{A/B}_{i,t}$ represents the proportion of trading that is accounted for by investors from Group A purchasing shares from investors in Group B. "Ind" and "Inst" denote individuals and institutions, respectively. Panel A reports results where the A/B trade variables are calculated as the proportion of total trade volume each day, and Panel B uses trade variables calculated as the proportion of daily turnover (shares traded divided by number of shares outstanding) each day. The regression is estimated using the Fama-MacBeth ("FM") approach in two ways: averaging results from time-series regressions by stock or from cross-sectional regressions by date. t-statistics for the FM by date regressions are calculated using standard errors robust to heteroscedasticity and autocorrelation using the Newey and West (1987) adjustment with five lags. The columns labeled "Sig. at 5%" report the percentage of regressions in which the coefficient is significantly negative or positive at the 5% level. Cross-sectional and time-series average standard deviations of the independent variables are reported in the columns labeled "C.S. Std. Dev" and "T.S. Std. Dev," respectively. Statistical significance at the 10%, 5% and 1% levels is denoted by †, *, and **, respectively.

		FM by stock					FM by date				
				Sig. at 5%		C.S.			Sig. at 5%		T.S.
		Estimate	t-stat	Neg.	Pos.	Std. Dev	Estimate	t-stat	Neg.	Pos.	Std. Dev
Panel A: Proportion of Daily Trading Volume											
Intercept		0.0022	3.72**	0%	44%		0.0011	3.16**	20%	26%	
$R_{i,t-1}$		-0.0634	-5.89**	52%	17%	0.0011	-0.0837	-11.46**	28%	11%	0.0009
Buyer:	*Seller:*										
Institutions	Institutions	0.0092	1.35	8%	11%	0.1481	0.0003	0.50	4%	4%	0.1691
	Households	0.0316	5.36**	0%	61%	0.1151	0.0189	16.09**	2%	11%	0.1180
Households	Institutions	-0.0423	-4.74**	72%	0%	0.1145	-0.0172	-12.70**	11%	3%	0.1130
	Households	-0.0025	-0.24	27%	0%	0.2002	-0.0023	-4.41**	9%	6%	0.2306
Panel B: Proportion of Daily Turnover											
Intercept		-0.0020	-8.45**	42%	3%		-0.0008	-3.06**	24%	19%	
$R_{i,t-1}$		-0.0791	-8.86**	58%	7%	0.0011	-0.1338	-19.87**	36%	7%	0.0009
Buyer:	*Seller:*										
Institutions	Institutions	0.0026	4.57**	0%	42%	0.9767	0.0011	3.69**	4%	10%	0.7566
	Households	0.0481	4.29**	1%	76%	0.5838	0.0390	19.62**	2%	27%	0.3728
Households	Institutions	-0.0339	-3.24**	48%	5%	0.2920	-0.0167	-10.15**	20%	6%	0.2782
	Households	0.0479	2.54*	2%	69%	0.3204	0.0033	2.27*	14%	18%	0.4697
Number of observations		134196					134196				
Number of FM regressions		106 (stocks)					2184 (days)				

Table 3.6: Returns and Group Interaction—VAR Results

This table presents coefficient estimates from the dynamic structural VAR(10) in equation (3.3),

$$\mathbf{A}_0 \mathbf{y}_{i,t} = \sum_{k=1}^{10} \mathbf{A}_i \mathbf{y}_{i,t-k} + \boldsymbol{\eta}_t,$$

where $\mathbf{y}_{i,t} = (\text{Ind/Ind}_{i,t}, \text{Inst/Inst}_{i,t}, \text{Ind/Inst}_{i,t}, \text{Inst/Ind}_{i,t}, R_{i,t})'$. The notation $\text{A/B}_{i,t}$ denotes the proportion of trading in stock i on date t that comes from Group A purchasing shares from Group B, and "Ind" and "Inst" denote individuals and institutions, respectively. $\mathbf{A}_0$ is the lower diagonal matrix from the triangular factorization of the error covariance in the reduced-form VAR (equation (3.2)). Separate regressions are estimated for each of the 106 stocks in the sample. The table reports cross-sectional averages of coefficient estimates for $k = 0, 1, 2$. Standard errors, calculated from the cross-sectional distribution of the coefficient estimates, are reported in parentheses. For brevity, the constant terms and lags of order three and higher are omitted. Statistical significance at the 10%, 5% and 1% levels is denoted by †, *, and **, respectively.

	Elements of $\mathbf{y}_{i,t-k}$				
Equation	Ind/Ind	Inst/Inst	Ind/Inst	Inst/Ind	Return
			$k = 0$		
Ind/Ind	–	–	–	–	–
Inst/Inst	-0.1303 (0.0355)**	–	–	–	–
Ind/Inst	0.0592 (0.0255)*	-0.0041 (0.0054)	–	–	–
Inst/Ind	0.0859 (0.0325)**	-0.0074 (0.0087)	-0.1091 (0.0057)**	–	–
Return	-0.0161 (0.0046)**	0.0059 (0.0041)	-0.0433 (0.0070)**	0.0344 (0.0064)**	–
			$k = 1$		
Ind/Ind	0.1489 (0.0074)**	-0.0201 (0.0093)*	0.0064 (0.0073)	0.0002 (0.0069)	-0.0135 (0.0213)
Inst/Inst	0.0227 (0.0200)	0.0964 (0.0057)**	0.0415 (0.0089)**	0.0488 (0.0063)**	-0.0303 (0.0174)†
Ind/Inst	0.0061 (0.0077)	0.0181 (0.0060)**	0.1493 (0.0062)**	-0.0083 (0.0039)*	-0.0560 (0.0130)**
Inst/Ind	0.0113 (0.0067)†	0.0339 (0.0094)**	-0.0045 (0.0035)	0.1451 (0.0116)**	0.0231 (0.0131)†
Return	0.0047 (0.0057)	0.0044 (0.0026)†	0.0019 (0.0024)	-0.0048 (0.0020)**	-0.0680 (0.0107)**
			$k = 2$		
Ind/Ind	0.0773 (0.0066)**	-0.0093 (0.0133)	0.0054 (0.0073)	0.0013 (0.0063)	0.0191 (0.0163)
Inst/Inst	-0.0157 (0.0242)	0.0521 (0.0044)**	0.0005 (0.0094)	0.0249 (0.0091)**	0.0107 (0.0157)
Ind/Inst	0.0119 (0.0051)**	0.0096 (0.0044)*	0.0760 (0.0057)**	-0.0103 (0.0037)**	-0.0516 (0.0132)**
Inst/Ind	0.0004 (0.0056)	0.0110 (0.0045)**	0.0056 (0.0039)	0.0696 (0.0050)**	0.0567 (0.0119)**
Return	0.0023 (0.0025)	0.0011 (0.0015)	0.0010 (0.0018)	-0.0006 (0.0030)	-0.0260 (0.0055)**

Table 3.7: Returns and Group Interaction—Weekly and Monthly

This table presents the results of the regression

$$R_{i,t} = \alpha + \beta R_{i,t-1} + \gamma_1 \text{Inst/Inst}_{i,t} + \gamma_2 \text{Inst/Ind}_{i,t} + \gamma_3 \text{Ind/Inst}_{i,t} + \gamma_4 \text{Ind/Ind}_{i,t} + \epsilon_{i,t},$$

where the notation $\text{A/B}_{i,t}$ represents the proportion of trading that is accounted for by investors from Group A purchasing shares from investors in Group B. "Ind" and "Inst" denote individuals and institutions, respectively. Panels A and B report results for regressions using data aggregated into weekly and monthly observations, respectively. The regression is estimated using the Fama-MacBeth ("FM") approach in two ways: averaging results from time-series regressions by stock (Columns 3 and 4), or from cross-sectional regressions by date (Columns 5 and 6). t-statistics for the FM by date regressions are calculated using standard errors robust to heteroscedasticity and autocorrelation using the Newey and West (1987) adjustment with five lags. Statistical significance at the 10%, 5% and 1% levels is denoted by †, *, and **, respectively.

		FM by stock		FM by date	
		Estimate	t-stat	Estimate	t-stat
		Panel A: Weekly Horizon			
Intercept		0.0057	5.43**	0.0056	2.77**
$R_{i,t-1}$		-0.0355	-3.82**	-0.0713	-5.61**
Buyer:	*Seller:*				
Financial institutions	Financial institutions	0.0320	1.14	0.0011	0.42
	Households	0.1882	3.87**	0.0630	8.37**
Households	Financial institutions	-0.2292	-5.59**	-0.0680	-11.07**
	Households	-0.0356	-0.42	-0.0103	-3.09**
Number of observations		33124		33124	
Number of FM regressions		106 (stocks)		456 (weeks)	
		Panel B: Monthly Horizon			
Intercept		0.0117	1.12	0.0290	2.63**
$R_{i,t-1}$		0.0001	0.01	-0.0103	-0.33
Buyer:	*Seller:*				
Financial institutions	Financial institutions	-0.0866	-1.13	-0.0197	-0.77
	Households	1.1009	5.01**	0.1689	3.12**
Households	Financial institutions	-1.7529	-2.97**	-0.2742	-8.52**
	Households	2.4445	1.73	-0.0295	-1.49
Number of observations		7680		7680	
Number of FM regressions		106 (stocks)		104 (months)	

Table 3.8: Returns and Group Interaction—Intraday Price Evidence

This table presents results from Fama-MacBeth regressions

$$\bar{b}^I_{i,t} - \bar{b}^H_{i,t} = \beta_1 \cdot \mathbf{1}_{\{R_{i,t}>0\}} + \beta_2 \cdot \mathbf{1}_{\{R_{i,t}<0\}} + \epsilon_{i,t}$$

and

$$\bar{s}^I_{i,t} - \bar{s}^H_{i,t} = \beta_1 \cdot \mathbf{1}_{\{R_{i,t}>0\}} + \beta_2 \cdot \mathbf{1}_{\{R_{i,t}<0\}} + \epsilon_{i,t},$$

where $\bar{b}$ and $\bar{s}$ denote the average purchase and sale prices, respectively, for stock i on date t, and superscript I and H denote institutions and households, respectively. The indicator function, $\mathbf{1}_{\{\cdot\}}$, takes a value of one only when the condition in curly brackets is true; otherwise it is zero. Results for these specifications are presented in Columns (1) and (3), respectively. Columns (2) and (4) show results for an alternative specification in which the price difference is allowed to vary with the magnitude of the return, $|R_{i,t}|$. For each stock and day, the average price at which individuals purchase (sell) shares is subtracted from the average price at which institutions purchase (sell) shares. This difference is used as one observation in a time-series regression for each stock. The table reports the mean coefficients estimates, with t-statistics (in parentheses) derived from the cross-sectional distribution of coefficient estimates. Statistical significance at the 10%, 5% and 1% levels is denoted by †, *, and **, respectively.

	Dependent Variable			
	$\bar{b}^I_{i,t} - \bar{b}^H_{i,t}$		$\bar{s}^I_{i,t} - \bar{s}^H_{i,t}$	
	(1)	(2)	(3)	(4)
$\mathbf{1}_{\{R_{i,t}>0\}}$	0.0025		-0.0011	
	(6.46)**		(-1.77)†	
$\mathbf{1}_{\{R_{i,t}<0\}}$	0.0002		-0.0016	
	(0.68)		(-6.24)**	
$\mathbf{1}_{\{R_{i,t}>0\}} \times \lvert R_{i,t}\rvert$		0.0649		-0.0005
		(11.67)**		(-0.09)
$\mathbf{1}_{\{R_{i,t}<0\}} \times \lvert R_{i,t}\rvert$		0.0026		-0.0718
		(0.37)		(-10.66)**
Number of observations	127501	127501	127501	127501
Number of regressions (stocks)	106	106	106	106

Table 3.9: Returns Following Trade

This table presents results for the regression

$$R_{i,t} = \alpha + \beta_1 \cdot \text{Ind/Ind}_{i,t-1} R_{i,t-1} + \beta_2 \cdot \text{Inst/Inst}_{i,t-1} R_{i,t-1} + \gamma \cdot \text{Controls} + \epsilon_{i,t},$$

for stock i on date t. "Inst" denotes financial institutions, and "Ind" denotes individual investors. Pairs of groups represent purchasing of shares by the first group from the second group. For example, "Ind/Inst" is the proportion of trading accounted for by individuals buying shares from financial institutions. The regression is estimated using the Fama-MacBeth ("FM") approach in two ways: averaging results from time-series regressions by stock (Columns 2 and 3), or from cross-sectional regressions by date (Columns 4 and 5). t-statistics for the FM by date regressions are calculated using standard errors robust to heteroscedasticity and autocorrelation using the Newey and West (1987) adjustment with five lags. Statistical significance at the 10%, 5% and 1% levels is denoted by †, *, and **, respectively.

	FM by stock		FM by date	
	Estimate	t-stat	Estimate	t-stat
$\text{Ind/Ind}_{i,t-1} \times R_{i,t-1}$	-0.2534	-4.09**	-0.1840	-3.50**
$\text{Inst/Inst}_{i,t-1} \times R_{i,t-1}$	0.0103	0.36	0.0443	1.26
Controls:				
Intercept	0.0011	5.41**	0.0009	2.52*
$R_{i,t-1}$	-0.0281	-2.63**	-0.0559	-6.62**
$\text{Ind/Ind}_{i,t-1}$	-0.0012	-1.31	-0.0009	-1.19
$\text{Inst/Inst}_{i,t-1}$	0.0006	1.18	0.0004	0.72
$\text{Ind/Inst}_{i,t-1}$	-0.0018	-2.55*	-0.0033	-5.84**
$\text{Inst/Ind}_{i,t-1}$	0.0022	3.09**	0.0013	2.37*
Number of observations	151324		151324	
Number FM of regressions	106 (stocks)		2184 (days)	

Table 3.10: Institutions' Response to Individual Trading

This table presents results for the regressions

$$\text{Inst/Ind}_{i,t} = \alpha + \beta_1 \cdot \text{Ind/Ind}_{i,t-1} R_{i,t-1} + \beta_2 \cdot \text{Inst/Inst}_{i,t-1} R_{i,t-1} + \gamma \cdot \text{Controls} + \epsilon_{i,t}$$

and

$$\text{Ind/Inst}_{i,t} = \alpha + \beta_1 \cdot \text{Ind/Ind}_{i,t-1} R_{i,t-1} + \beta_2 \cdot \text{Inst/Inst}_{i,t-1} R_{i,t-1} + \gamma \cdot \text{Controls} + \epsilon_{i,t},$$

for stock i on date t. "Inst" denotes financial institutions, and "Ind" denotes individual investors. Pairs of groups represent purchasing of shares by the first group from the second group. For example, "Ind/Inst" is the proportion of trading accounted for by individuals buying shares from financial institutions. The dependent variable in Panel A is the proportion of trading that comes from individuals purchasing shares from institutions (Ind/Inst$_{i,t}$), and in Panel B it is the proportion of trading that comes from institutions purchasing shares from individuals (Inst/Ind$_{i,t}$). The regression is estimated using the Fama-MacBeth ("FM") approach in two ways: averaging results from time-series regressions by stock (Columns 2 and 3), or from cross-sectional regressions by date (Columns 4 and 5). t-statistics for the FM by date regressions are calculated using standard errors robust to heteroscedasticity and autocorrelation using the Newey and West (1987) adjustment with five lags. Statistical significance at the 10%, 5% and 1% levels is denoted by †, *, and **, respectively.

	FM by stock		FM by date	
	Estimate	t-stat	Estimate	t-stat
Panel A: Dependent Variable=Inst/Ind$_{i,t}$				
Ind/Ind$_{i,t-1} \times R_{i,t-1}$	-0.6939	-3.12**	-1.7931	-4.02**
Inst/Inst$_{i,t-1} \times R_{i,t-1}$	0.1050	1.20	-0.0723	-0.43
Controls:				
Intercept	0.1240	21.45**	0.1020	48.13**
$R_{i,t-1}$	0.0557	1.67†	0.1645	2.34*
Ind/Ind$_{i,t-1}$	0.0239	2.81**	0.0519	7.10**
Inst/Inst$_{i,t-1}$	0.0016	0.44	-0.0241	-7.28**
Ind/Inst$_{i,t-1}$	-0.0643	-11.14**	-0.0180	-3.88**
Inst/Ind$_{i,t-1}$	0.2180	22.09**	0.2780	47.83**
Panel B: Dependent Variable=Ind/Inst$_{i,t}$				
Ind/Ind$_{i,t-1} \times R_{i,t-1}$	0.3886	2.03*	0.9371	2.42*
Inst/Inst$_{i,t-1} \times R_{i,t-1}$	-0.1958	-2.38*	-0.1951	-1.54
Controls:				
Intercept	0.1186	19.21**	0.0847	42.22**
$R_{i,t-1}$	0.0308	0.89	-0.0116	-0.20
Ind/Ind$_{i,t-1}$	0.0416	5.55**	0.0984	14.26**
Inst/Inst$_{i,t-1}$	0.0106	2.88**	-0.0150	-4.57**
Ind/Inst$_{i,t-1}$	0.2452	24.36**	0.3032	61.63**
Inst/Ind$_{i,t-1}$	-0.0483	-10.48**	0.0005	0.13
Number of observations	151324		151324	
Number FM of regression	106 (stocks)		2184 (days)	

Figure 3.1: Stylized Timeline of Price Path Around Trade

This figure shows alternative price paths following a trade. The trade takes place at time t_0. In the top two figures, the trade is initiated by the buyer, and the price immediately increases. In the bottom two figures, the trade is initiated by the seller, and the price immediately decreases. If the trade initiator is uninformed, prices subsequently revert. If the trade initiator is informed, there is no such reversion.

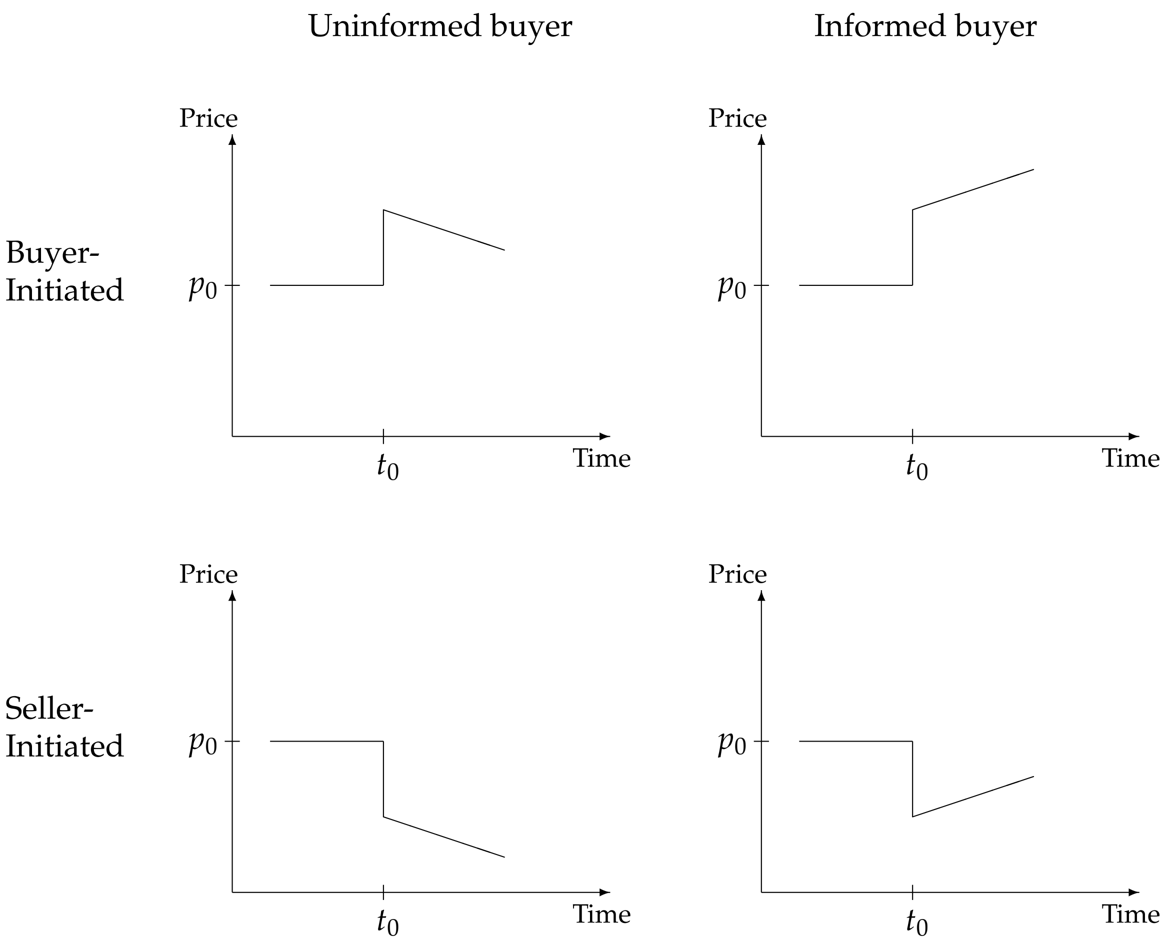

Figure 3.2: Cumulative Price Impact Functions

This figure plots the accumulated orthogonalized impulse response function for the structural dynamic VAR in equation (3.3),

$$\mathbf{A}_0 \mathbf{y}_{i,t} = \sum_{k=1}^{10} \mathbf{A}_i \mathbf{y}_{i,t-k} + \boldsymbol{\eta}_t,$$

where $\mathbf{y} =$ (Ind/Ind, Inst/Inst, Ind/Inst, Inst/Ind, Return)$'$ and $\mathbf{A}_0$ is the lower diagonal matrix from the triangular factorization of the error covariance in the reduced-form VAR (equation (3.2)). The graphs show the cumulative effect on returns (in basis points) of a one standard deviation shock ($\eta_{i,t}$) to each of the trade variables during the trading day that occurs between time -1 and time 0. Returns are calculated using closing prices each day. Inst/Ind (solid) is the proportion of trading accounted for by institutions purchasing shares from individuals; Ind/Ind (long dash) represents trading between two individuals; Inst/Inst (short dash) represents trading between two institutions; and Ind/Inst (dash-dot) represents individuals purchasing shares from institutions.

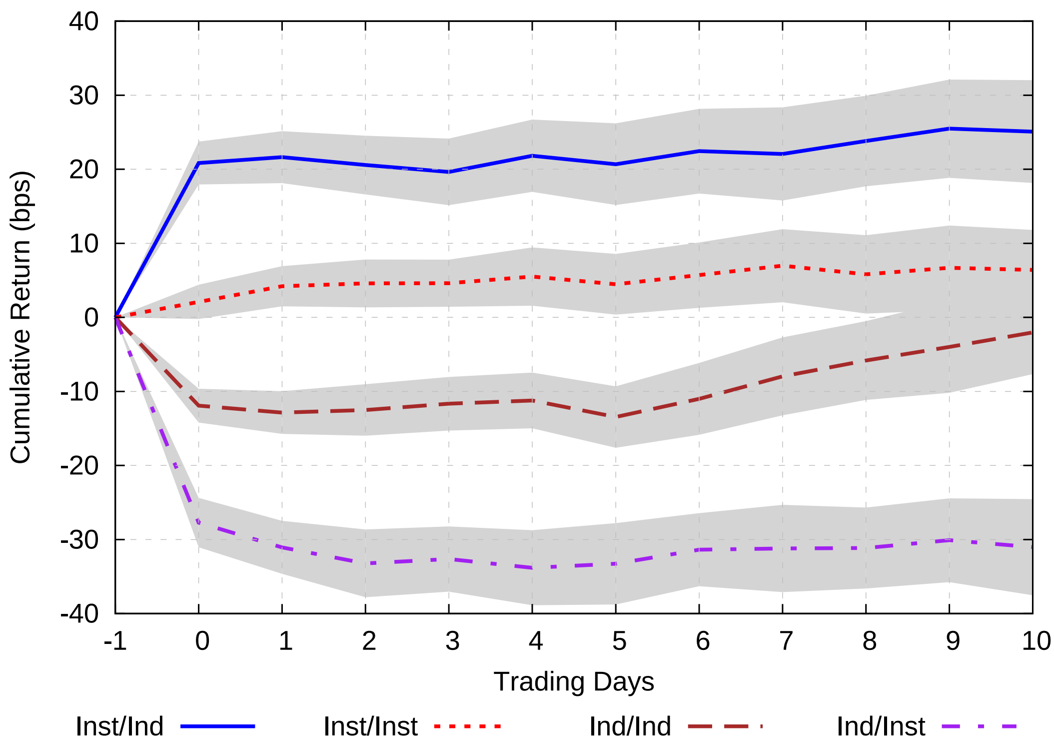

Chapter IV

When Are Individual Investors Informed?

The decisions of individuals lie at the foundation of economic theory. Unfortunately, there are very few real-world settings where these decisions and their economic consequences are directly observed at the individual level. One setting that has proved fruitful for examining such decisions and outcomes is financial markets. In a series of seminal papers, Odean (1998, 1999) and Barber and Odean (2000, 2001) examined the trading decisions of a large number of investors at a large discount brokerage firm in the U.S. Subsequently, a large literature has developed analyzing the behavior of individual investors in a number of settings.[1]

The main stylized fact that emerges from this literature is that individuals make poor investment decisions. For example, individuals trade too much, generating transaction costs without increasing their returns; they make bad choices, as stocks they sell subsequently earn higher returns than stocks they buy; and they

[1]Some representative papers from this literature include Grinblatt and Keloharju (2000, 2001a, 2001b), Benartzi (2001), Goetzmann and Kumar (2008), Ivković and Weisbenner (2005), Ivković, Sialm, and Weisbenner (2006), and Feng and Seasholes (2005).

hold on to their losing stocks but sell their winning stocks, which is a sure way to lose money in the presence of momentum.

In this chapter, I revisit the question of whether individual investors make good investment decisions. I find that the results of previous papers are correct, but incomplete. In particular, I begin by focusing on one of the earlier findings of this literature: that post-purchase returns are lower than post-sale returns (Odean 1999). I show that while this is indeed the case unconditionally, it is not the case when we consider those investment decisions that are more likely, *ex ante*, to be informed.

Broadly speaking, there are two reasons an investor might buy or sell a stock: liquidity, or information. As an example of a liquidity-motivated trade, an investor may have earned income that she wishes to save, and chooses to invest in equities. Or she might have unexpected costs, and chooses to sell some securities to generate the needed cash. In both of these examples, the investor either purchases one or more stocks *or* sells one or more stocks; she does not transfer money from one stock to another. Suppose instead that the investor chooses to sell one stock she owns and use the proceeds to purchase another stock. In this case, she may have received information that causes her to believe the asset she owns is going to decline in value, or that the stock she chooses to buy is going to increase in value; in either case, it is clear that she believes the stock she buys will earn a higher return than the stock she sells. That is, she is forced to place an "extra" trade relative to what she would have done if she were not cash-constrained.[2]

The economic intuition that comes from these examples suggests a way to identify which trades are more likely to be informed. If we observe an investor liquidating their holdings in one stock and then using the proceeds to purchase another stock, it is more likely that these transactions are motivated by information than other transactions. I use this intuition to identify informed trades among individual investors. When an investor sells shares in one stock and subsequently purchases shares in another stock, within a short period of time, I call this a "substitution" trade. Obviously, not all of these trades are perfectly informed, but the results below provide strong evidence that this classification does capture a meaningful distinction between informed and uninformed trades. The details of the classification procedure I employ are provided in section 4.3, below.

I show that the post-purchase and post-sale returns to substitution trades are significantly better than the returns for non-substitution ("regular") trades. This finding suggests that investors are not as bad at making investment decisions as previously documented. In particular, it is possible to identify, *ex ante*, a subset of investment decisions that are likely to be driven by information. In this regard, this chapter is related to the work of Coval, Hirshleifer, and Shumway (2005), who find that a subset of investors are able to beat the market consistently, and Seru,

[2] Alternatively, she may not believe that the stock will earn a higher return than stocks she already holds, but wishes to rebalance her portfolio to change the riskiness of her holdings. The extent to which this is true will make it more difficult for me to find the results I report below.

Shumway, and Stoffman (2007), who find that investors learn to avoid behavioral biases as they become experienced.

The present study is also related to the recent work of Alexander, Cici, and Gibson (2007), who find that liquidity-motivated investment choices of mutual fund managers earn lower returns than trades that are more likely to be motivated by private information. The authors distinguish between transactions necessitated by investor fund flows and those that are not, and use a similar argument as I do here to claim that the latter trades are more likely to be informed. For example, they note that a fund manager who purchases a stock even when investors are withdrawing money from the mutual fund is likely to believe strongly that the stock is undervalued, whereas if the purchase is made in the presence of fund inflows, it is more likely to be just a reallocation of holdings from excess cash to stocks. They find that managers do, in fact, beat the market on these trades, and not when they trade because of investor inflows and outflows.

I also show that the pattern of returns before and after regular trading differs from that surrounding substitution trades. When substitution trades are dominated by purchases, prices rise and do not subsequently revert. In contrast, prices respond negatively to regular purchases, and subsequently revert. This provides additional evidence that there is distinctly different information content in substitution and non-substitution trades. These results complement the recent work of Barber, Odean, and Zhu (2006), Hvidkjaer (2006), and Kaniel, Saar, and Titman (2008), all of whom examine the relation between individual investor order flow and stock returns.

A particular strength of this chapter is that I document consistent results using two large datasets from different countries over different time periods. This adds considerable support the notion that the finding reflects a result of general economic behavior, and not a quirk of one particular setting, or a statistical fluke.

The remainder of this chapter is organized as follows. In Section 4.2, I discuss the main datasets used in this study. I outline my methods and present results in Section 4.3. Section 4.4 concludes.

4.1 Related literature

The present study contributes to the growing literature on the relation between the trading of individual investors and stock returns.

Obizhaeva (2007) examines the price impact of trades by institutions that are transitioning from existing positions to a target portfolio. These transitions frequently occur after a change in the portfolio manager, and the author provides evidence that transition purchases are informed while sales are likely to be liquidity motivated.

Examining quarterly 13(f) filings of institutional stock holdings, Sias, Starks, and Titman (2006) provide evidence that trading by institutions causes the observed correlation between quarterly changes in institutional ownership and returns. In addition, they find evidence in support of the notion that institutional investors are better informed than other traders, and that their information is incorporated into prices as they trade.

San (2007) also uses 13(f) filings to conclude that individuals earn higher returns than institutions during the period 1981 to 2004. A potential problem with this approach, however, is the assumption that traders who do not file 13(f) reports are individuals, while in fact this group would include hedge funds and smaller money managers.[3] Nevertheless, this is the same approach taken by Gompers and Metrick (2001), who find that aggregate institutional performance is better than individual performance by 0.67% per year, and that this difference is driven more by demand shocks caused by institutional preferences for liquid stocks than by institutions being "smarter" than individuals.

Kaniel, Saar, and Titman (2008) analyze trades placed by individuals in NYSE stocks during 2000–2003. They find that individuals are contrarian at a monthly horizon: they purchase stocks after recent declines and sell after recent increases. Moreover, they document positive excess returns in months following high levels of individual buying, which is consistent with individuals supplying liquidity to institutions.

4.2 Data

In addition to the Finish data outlined in Chapter I, I use a sample of trading records from investors at a large discount brokerage firm in the United States.[4]

[3]The Securities Exchange Act requires institutional investment managers who exercise investment discretion over securities having an aggregate fair market value on the last trading day of any month of at least $100 million to file a 13(f) report within 45 days of the end of each calendar quarter. These institutions are required to report only those positions of greater than 10,000 shares or $200,000 in market value. See `http://www.sec.gov/about/forms/form13f.pdf`.

[4]I thank Terry Odean for sharing these data.

This data set contains 3.1 million observations, covering the trades of 126,488 accounts from January, 1991 to November, 1996. Each observation consists of an account identifier, a security identifier (CUSIP), trade date, number and price of shares traded and a buy/sell indicator.

I use the CUSIP to merge the trade data with data from the Center for Research in Security Prices (CRSP). I extract the return series, `ret`, to calculate post-transaction returns, as discussed below. This variable includes both the capital gains and any dividends or other payouts earned by shareholders.

4.3 Methods and Results

In this section I discuss my methods, and present the results. I first examine the general pattern of returns following purchases and sales by individual investors, and confirm the results of Odean (1999) in different datasets. These results are "unconditional" in the sense that I am not classifying the trades in any way before calculating post-transaction returns. I then explain how I identify those trades that are more likely to be informed ("substitution" trades), and show that the pattern of returns for these trades is, in fact, quite different. While *unconditionally* the returns following a purchase are lower than the returns following a sale, substitution purchases earn higher returns and substitution sales earn lower returns than non-substitution ("regular") trades. This provides significant evidence that substitution trades are more informed than other trades.

4.3.1 Unconditional post-transaction returns

A straightforward way to assess the success of an individual's trading record is to examine the return of stocks purchased and sold subsequent to the transaction date. Therefore, for each transaction, I calculate the post-transaction holding period return, $R_{i,j,h}$, for investor i's jth transaction over a horizon of h trading days. Holding period returns are calculated as the percentage increase in a cumulative return index that includes dividend payments. To ensure that my results are not affected by bid-ask bounce, I calculate returns using the closing price on the day after the transaction.

These post-transaction returns in the U.S. data are summarized in Table 4.1 for purchases and sales, separately. Odean's (1999) finding that post-sale returns are higher than post-purchase returns is confirmed in these data, which partially overlap with the data used in that study. For example, in the 63 trading days (3 months) following a trade, average returns are 3.34% for purchases and 3.88% for sales, for a 54 basis point difference. As Odean points out, this implies that investors would have earned higher returns (approximately 2% in this case) if they had simply abstained from trading. It is also interesting to note in this sample that individuals were approximately 20% more likely to purchase shares than to sell shares in this sample.

The results in Table 4.1 are calculated using the same approach as Odean (1999). That is, returns are aggregated across all accounts without first aggregating by each account. Therefore, the post-purchase returns of one investor are being compared to the post-sale returns of another investor. If subsets of investors have different skill levels, and are present in the data over different time periods, this approach could lead to incorrect inference. In order to rule this out, I repeat the analysis using two different approaches.

First, I run the regression

$$R_{i,j} = \alpha_i + \beta \cdot 1_{\{j=\text{Buy}\}} + \epsilon_{i,j}, \tag{4.1}$$

where the α_i are individual fixed effect and $1_{\{\cdot\}}$ is a dummy variable. The regression is repeated for each post-transaction return horizon. The individual-specific intercept controls for all time-invariant individual-specific heterogeneity, such as different levels of skill. Results from these regressions are shown in Table 4.2. Accounting for individual-specific heterogeneity mitigates the differences in returns, although purchases still significantly underperform sales. For example, at the 84-day horizon, buys underperform sells by 45 basis points, compared to the 56 basis points calculated above. Nevertheless, the results remain strongly significant at conventional levels.

Second, I calculate average buy- and sell-returns for each account, and then calculate the difference in means for each account before taking the cross-sectional mean across all accounts. This approach is designed to examine whether any particular investor's purchases can be expected to underperform that same investor's sales. The results from this approach are shown in Table 4.3, and are again quite similar to those already reported. Means in Panel A are equally weighted, and in Panel B are weighted by the number of trades made by the account. Panel C reports equally weighted means excluding all accounts that placed fewer than ten trades. With the exception of the 5-day return in Panel C, all results are significantly negative at the 1% level. The weighted average results in Panel B are probably the most conservative, and are very close to the coefficient estimates in Table 4.2.

The results presented in this subsection confirm the findings of Odean (1999). These results are robust to using various statistical approaches, and are confirmed in the Finnish data set.[5] In the rest of this chapter, I ask whether these results continue to hold when we differentiate between trades that are likely to be informed ("substitution" trades) and those that aren't.

4.3.2 Substitution trades

As outlined above, I define a "substitution trade" as a group of transactions comprising at least two different stocks, where a purchase of one stock is followed by a sale of another within two to five trading days. Moreover, the total value of shares purchased must be "close" to the total value of shares sold, where "close" is defined as within two to 25 percent of the value traded.[6] These are trades where the investor has made a clear decision to substitute one security for another.

For example, suppose we observe the following transactions for an investor:

[5]In the interests of space, the various sets of tests for Finnish data are not reported in this section, but the fact that they are similar to the U.S. results will become clear in the remainder of this chapter. All results are available from the author upon request.

[6]The results presented below are robust to using a longer window of up to 15 trading days.

Transaction	Date	Stock ID	Buy/Sell	Value Traded
1	1/15/02	A	Buy	$800
2	2/12/02	A	Sell	600
3	2/13/02	B	Buy	585
4	2/18/02	C	Buy	500
5	5/21/02	B	Sell	700

In this case, transactions 1, 4, and 5 would be classified as "regular" trades, while transactions 2 and 3 would be classified as one "substitution" trade.

The number of transactions classified as being part of a substitution trade is reported in Table 4.4 for the U.S. data. The criteria for classification become less selective as we increase the number of days in the window, or how close the value of shares purchased and sold is (denoted "Cutoff" in the table), so the number of substitution trades increases as we move down and to the right in the table. If we look only at transactions that occur within one of each other and where the total value of shares purchased is within 1% of shares sold, 2.9% of transactions are classified as substitution trades. At the other end of the spectrum, 33.1% of transactions are classified as substitution trades when we allow transactions to occur over a 15 day period, and having a purchase value within 25% of the value of shares sold.

It is difficult to gauge how reasonable these numbers are, although some guidance is available from the existing literature. Using a structural model similar to that of Glosten and Milgrom (1985), Easley, Kiefer, O'Hara, and Paperman (1996) estimate the probability of informed trading in a number of large NYSE-traded stocks. Their estimation procedure uses only data on the total number of buyer- and seller-initiated transactions in a stock on each day, and therefore includes the trades institutions, which are likely to make up the bulk of trading in most large stocks. They obtain estimates on the order of 15% of trading being informed. Since I am looking exclusively at trading by individual investors, it seems reasonable to expect a lower percentage of informed trades in my sample.

For the remainder of this chapter, I use the 2-day/5% criteria for classification. That is, a trade is classified as a substitution trade if there are no more than two days between the first transaction and the last transaction, and if the total value of shares purchased does not differ from the total value of shares sold by more than 5%. While this choice is somewhat arbitrary, I confirm in unreported tests that the results are not sensitive to this particular choice of parameters. The general pattern of these unreported results is that they are weaker as we lengthen the window of classification and increase the difference allowed between purchase and sale amounts. This makes sense, since more and more trades are likely to be incorrectly classified as information driven when we use these looser criteria.

4.3.3 Returns to substitution trades

Results presented in the existing literature, and confirmed above, show that investors on average make the wrong trading decisions: they purchase stocks that subsequently decrease in value, and sell stocks whose prices subsequently increase. In this section, I investigate whether this pattern continues to hold up when we consider those trades that are more likely to be informed (that is, the "substitution" trades). Results from the U.S. data are presented first, and then the results from Finland.

The key result of this chapter comes from regression tests of post-transaction returns. In particular, I report below coefficient estimates from the regression

$$R_{i,t} = \alpha + \beta_1 \cdot \text{Buy} + \beta_2 \cdot \text{Substitution} + \beta_3 \cdot \text{Buy} \times \text{Substitution} + \epsilon_{i,t}, \qquad (4.2)$$

where R denotes returns, Buy is a dummy variable that takes a value of one for purchases and zero for sales, and Substitution is a dummy variable that takes a value of one when a trade is classified as part of a substitution trade, and zero otherwise. 56% of all transactions are purchases.

The interpretation of the regression coefficients is made particularly easy by the fact that all regressors are dummy variables. As such, the coefficients correspond to changes in the mean value of the dependent variable for different constellations of the dummy variables. The average post-sale return for a trade that is not classified as a substitution trade (a "regular" trade) is given by α. The β_1 coefficient corresponds to the difference between a regular purchase and a regular sale. The average post-sale return for a substitution trade is $\alpha + \beta_2$, whereas the post-purchase return to a substitution trade is $\alpha + \beta_1 + \beta_2 + \beta_3$. Note as well that this difference-in-difference specification implies that the coefficient on the interaction term corresponds to the difference in the difference between post-sale and post-purchase returns for regular trades and substitution trades:

$$\beta_3 = [E(R_{i,t} \mid \text{Buy, Substitution}) - E(R_{i,t} \mid \text{Sell, Substitution})] \\ - [E(R_{i,t} \mid \text{Buy, Regular}) - E(R_{i,t} \mid \text{Sell, Regular})] \, .$$

A test of $\beta_3 = 0$ is therefore a test of Hypothesis 1, namely that substitution trades perform better than regular trades.

U.S. data

Panel A of Table 4.5 shows the results of four separate regressions as in (4.2) using post-transaction return horizons of 1, 3, 6, and 12 months, respectively, with the U.S. data. Importantly, standard errors are calculated using the robust "sandwich" estimator with clustering at the stock or fund level. This approach allows the return of a stock to be correlated across individuals, which makes sense because many individuals may hold the same stock at the same time. Using clustered standard errors drastically cuts the effective degrees of freedom in the model, and may be overly conservative, since it seems reasonable to argue that each individual investor's decision to transact is the appropriate unit of observation in this setting.

The coefficient on the Buy dummy is reliably negative, which again confirms that post-purchase returns are lower than post-sale returns. At a one year horizon, a stock increases on average 16.32% following a sale, and 14.3% following a purchase (2.02% less). For all horizons, these differences are highly statistically significant.

As noted above, the most relevant estimate is the coefficient on the interaction term, Buy $\times$ Substitution. Even with the clustered standard errors, this coefficient is positive and highly statistically significant at each horizon, indicating that the performance of purchases relative to sales is significantly improved for substitution trades compared to regular trades. Moreover, at each horizon the interaction coefficient is larger than the coefficient on the Buy dummy, so the worse performance of purchases than sales for regular trades is actually reversed for substitution trades. This can be seen in the row labeled $\beta_1 + \beta_3$: for example, at the one year horizon, while a regular buy earns on average 2.02% *less* than a regular sell, a substitution buy earns 39 basis points *more* than a substitution sell. The p-values for the F-statistics associated with these coefficient sums are reported in square brackets. Except for at the 3-month horizon, these results are not significantly greater than zero, but a one-sided test against a null hypothesis of $\beta_1 + \beta_3 < 0$ is strongly rejected at any conventional significance level.

It is interesting to note as well where the improvement of post-transaction returns is coming from for the substitution trades. The Substitution coefficient is negative, so substitution sales earn lower returns than regular sales; this is consistent with the notion that substitution trades are made by investors with information about which stocks to sell. The difference between a Substitution Buy and a Regular Buy is given in the row labeled $\beta_2 + \beta_3$. Here the statistical evidence is much weaker, although it appears that the results are somewhat stronger at the longer horizon. At the 1- and 3-month horizon substitution purchases earn no more than 9 basis points less than regular purchases, and at the 6- and 12-month horizons substitution buys earn 26 and 66 basis points more than regular buys.

In other words, the improvement in post-transaction returns for substitution trades is mainly due to sales, especially at shorter horizons. Investors placing substitution trades appear to be significantly better than investors placing regular trades at identifying which stocks to sell, but only somewhat better at choosing which stocks to purchase. One potential explanation for this asymmetry is that identifying a stock to sell from the relatively small number stocks currently held is easier than identifying a stock to purchase from the much larger set of stocks that the investor does not currently hold. Another possibility is that short sale constraints prevent negative information from being quickly impounded in stock prices, so individual investors can receive negative information about a stock they currently hold that isn't yet reflected in prices. Positive information, on the other hand, might be impounded in prices more quickly, giving individual investors less time to take advantage of information they may receive.

To ensure that the results presented so far are not driven by especially illiquid stocks, Panel B of Table 4.5 presents estimates from the same regressions, but the sample is restricted to just those stocks whose price is above $5. The results here are qualitatively the same as those in Panel A.

Focusing on stocks with prices above $5 also allows us to investigate a little more the possibility that short sale constraints make it easier for individual investors to identify stocks to sell than stocks to buy. Higher-priced stocks are somewhat less likely to have binding short sale constraints, although perhaps only marginally. Nevertheless, comparing Panel B to Panel A, post-purchase returns for Substitution buys are higher in the more liquid stocks of Panel B, and post-sale returns are lower for the more liquid stocks. For example, at the 12-month horizon, the difference between Regular sales and Substitution sales is 175 basis points using the entire sample, but 130 basis points for the higher-price sample; and the difference between substitution buys and regular buys is 66 basis points for the entire sample, but 82 basis points for the restricted sample. This is consistent with the short sale constraint story: for the stocks in Panel B, which are less likely to have binding short-sale constraints, investors are not as good at choosing which stocks to sell but better at choosing which stocks to buy than the entire sample considered in Panel A.

Finland data

Turning to the Finnish dataset, Table 4.6 presents results from the same type of post-transaction return regressions as those just presented. The results in Finland are remarkably consistent with those in the U.S.

In particular, there is a strong negative effect of the Buy dummy, so for non-substitution trades, purchases are followed by lower returns than sales. Again, the effect is reversed for Substitution trades: the interaction coefficient is highly significant and economically important. For example, at a one-year horizon, purchases are followed by 13.78% *lower* returns than sales for Regular trades, but Substitution purchases are followed by 31 basis points *higher* returns than Substitution sales, as shown in the row labeled $\beta_1 + \beta_3$. The difference for regular trades is strongly statistically significant, and statistically insignificant for substitution trades. And as in the U.S. data, the improvement in returns for substitution trades in Finland comes primarily from the sell side: Substitution purchases are followed by somewhat higher returns than regular purchases, although the difference is insignificant, but substitution sales are followed by much lower returns than regular sales, as can be seen by examining the substitution dummy coefficient and the sum of $\beta_2 + \beta_3$.

The results presented in this section provide strong evidence for the notion that Substitution trades are motivated by information much more than other trades. The evidence is remarkably consistent, especially considering the data are from substantially different periods (1991–1996 in the U.S. and 1995-2003 in Finland), and the substantially different structure of the two markets (approximately 2 million observations and 11,000 stocks in the U.S. and 6 million observations and 149 stocks in Finland).

4.3.4 Returns and Aggregate Buy Ratios

If substitution trades are more likely to be informed, we would expect them to display different time series properties than other trades. For example, if these trades contain private information we would expect them to have a permanent price impact. We would also expect these trades not to be related to past returns in a way that uninformed trades may be (if, for example, investors without information respond to price momentum). I investigate these questions in this section.

I begin by calculating an aggregate Buy Ratio for substitution trades and non-substitution trades. For each stock and each day (or week), I count the number of substitution buys and substitution sells. The substitution Buy Ratio is the number of substitution buys divided by the total number of substitution sells. Similarly, the buy ratio for regular trades is calculated as the total number of regular purchases divided by the total number of regular trades. I then estimate the regression

$$R_{i,t+h} = \alpha + \beta \cdot BR_{i,t} + \epsilon_{i,t}, \quad \text{for } h = -4, \ldots, 4, \tag{4.3}$$

where $BR_{i,t}$ denotes the Buy Ratio for stock i at time t. The substitution Buy Ratio or the regular Buy Ratio are used in separate regressions.

Estimated regression coefficients for β are reported in Table 4.7. Each cell presents the results of a separate regression. Results using a one-day aggregation window are reported in Panel A, and those using one week are reported in Panel B. Consider first the estimates in the column labeled "0" in Panel A. There is no significant relation between the regular buy ratio and contemporaneous returns. In contrast, there is a strong relation between the substitution buy ratio and contemporaneous returns. That is, the higher is the substitution buy ratio, the more we expect prices to increase; no such relation exists for the regular buy ratio. This is again consistent with the notion that substitution trades are more likely to be informed; informed purchases push up prices, while uninformed purchases have no effect.

Further evidence in support of the claim that substitution trades are more informed than regular trades is provided by the next four columns of the table. In the days following a high regular buy ratio, returns are significantly negative, suggesting price movements on "day 0" are quickly reversed. No such reversals are apparent in the days following a high substitution buy ratio, suggesting that price impact on day 0 is permanent. Looking at the four days preceding the buy ratio calculation, there is no evidence that the level of substitution buying or selling is related to previous returns. There is, however, some evidence that regular purchases are more likely to follow negative return days. This is borne out in the weekly data, which we turn to next.

Panel B shows coefficient estimates for the same types of regressions, although the buy ratios are now calculated using all regular or substitution trades within a week (Monday to Friday). The significant differences between the return response of stocks to the two types of buy ratios shown in Panel A are again found in Panel B. At the weekly horizon, there is still a highly significant price response to the contemporaneous substitution buy ratio, but also a highly significant price response in the opposite direction to the regular buy ratio (column "0"). In other words, weeks in which substitution trades are predominantly purchases are weeks with high returns, and weeks in which regular trades are dominated by purchases are weeks with low returns. Moreover, when returns have been low in the past four weeks, the regular buy ratio is higher, indicating contrarian trading behavior among regular trades. Substitution trading, on the other hand, is momentum trading in the sense that it is higher when the previous week's return is higher.

And as in the daily results, there is evidence of a price reversal following regular trading, but no such reversal following substitution trades.

The pattern of results in Table 4.7 provides strong evidence that substitution trades are fundamentally different from other trades. In particular, the results are consistent with the substitution trades being informed and regular trades being uninformed. We would expect informed trading to cause permanent price movements, and uninformed trading to be followed by price reversals. This is precisely what we see for the substitution and regular trades. We might also expect uninformed trading to respond to previous price movements, such as if investors purchase shares after a price decline because these are interpreted as a "buying opportunity." We would not expect strong effects of this sort for informed trading, and again, this is what we find for the regular and substitution trades, respectively.

4.4 Conclusion

This chapter presents a simple method for classifying the trades of individual investors "informed" or "uninformed." Starting with the economic argument that liquidity trading is likely to show up as purchases or sales of stocks and not moving money from one stock to another, I claim that trades are more likely to be informed when they involve selling shares of one stock and using the proceeds to buy a similar dollar value of another stock in a short period of time.

I implement this classification procedure in two large datasets from the U.S. and Finland, covering 1991–1996 and 1995–2003, respectively. The datasets are quite different; for example, investors in the U.S. sample trade upwards of 11,000 different securities, while investors in the Finland data trade 149 securities. Despite these differences, a remarkably consistent result emerges: substitution trades do appear to be informed. The post-transaction returns for substitution trades are significantly better than the post-transaction returns for non-substitution trades. Moreover, when substitution trades are dominated by purchases, returns are significantly positive, indicating a price impact. This price impact is not subsequently reversed for substitution trades, whereas there is a significant reversal for non-substitution trades.

While the results in this chapter provide substantial evidence that some individual investor trading is informed, they nevertheless represent a first-pass at identifying informed trading among individual investors. Surely, more sophisticated classification methods using more of the available data are possible. Such classifications would give allow a better understanding of the interaction between informed and uninformed investors in financial markets, and have the potential to greatly improve upon existing measures such as Easley, Kiefer, O'Hara, and Paperman's (1996) "PIN" measure. These and other issues are the focus of continuing research.

Table 4.1: Post-Transaction Returns in U.S. Data—Summary Statistics

This table presents summary statistics for returns subsequent to purchases and sales. The unit of observation is a transaction, and N denotes the total number of observations. Returns are calculated as the increase in a cumulative return index (including all payouts) beginning the day after a transaction, and over the horizon specified in each column. 5 days corresponds to one week, 21 days to one month, and 252, 504, and 756 trading days correspond to one, two, and three years, respectively.

	Horizon (in trading days)							
	5	21	63	84	126	252	504	756
	Purchases							
Mean	0.0030	0.0129	0.0334	0.0440	0.0646	0.1516	0.3650	0.5834
Std. Dev.	0.0812	0.1454	0.2490	0.2852	0.3540	0.5432	0.9445	1.2487
25th Pctl.	-0.0295	-0.0591	-0.1012	-0.1138	-0.1350	-0.1592	-0.1610	-0.1142
50th Pctl.	0.0000	0.0072	0.0214	0.0256	0.0356	0.0863	0.1951	0.3132
75th Pctl.	0.0315	0.0762	0.1434	0.1709	0.2162	0.3464	0.6267	0.9131
N	821,835	810,397	783,176	770,705	734,680	628,436	467,639	344,860
	Sales							
Mean	0.0036	0.0156	0.0388	0.0496	0.0736	0.1758	0.3954	0.6327
Std. Dev.	0.0675	0.1410	0.2392	0.2757	0.3469	0.5373	0.9081	1.2798
25th Pctl.	-0.0268	-0.0534	-0.0917	-0.1037	-0.1175	-0.1259	-0.1182	-0.0689
50th Pctl.	0.0000	0.0091	0.0256	0.0308	0.0437	0.1099	0.2269	0.3543
75th Pctl.	0.0298	0.0750	0.1456	0.1724	0.2180	0.3691	0.6715	0.9574
N	678,394	666,379	637,962	625,747	596,415	507,254	371,564	271,042

Table 4.2: Post-Transaction Returns in U.S. Data—Regression Tests

This table presents regression results for the fixed-effects regression

$$r_{i,j} = \alpha_i + \delta_m \cdot 1_{\{t \in m\}} + \beta \cdot 1_{\{j=\text{Buy}\}} + \epsilon_{i,j},$$

where α_i is the individual-specific intercept, and $1_{\{\cdot\}}$ is an indicator function. t denotes the date of the transaction, and $m \in \{1, \ldots, 70\}$ is a month index number. Post-purchase returns are lower than post-sale returns if $\beta < 0$. The number of observations used in the regression is denoted N Obs.

Horizon	β	t-statistic	N Obs.
5	0.0005	2.58	1,500,225
21	0.0006	1.99	1,476,772
63	-0.0024	-4.46	1,421,136
84	-0.0032	-5.03	1,396,450
126	-0.0059	-7.21	1,331,093
252	-0.0145	-10.28	1,135,688
504	-0.0108	-3.72	839,202
756	-0.0239	-5.24	615,901

Table 4.3: Post-Transaction Returns in U.S. Data—Cross-sectional Means

This table presents the cross-sectional means of an individual's average post-purchase return minus that individual's average post-sale return. Horizons are measured in trading days, with 5 days corresponding to one week, and 756 days to approximately three years. Num. Accts is the number of unique accounts used in the cross-sectional calculation. In Panel A, a simple average is calculated across all accounts. Panel B restricts the analysis only to those accounts that place at least 10 trades. Except where otherwise noted, all means are statistically different from zero at the 1% level.

Horizon	Num. Accts	Mean	Std. Dev.	*t*-statistic
Panel A: All Accounts				
5	67,862	-0.0007	0.0599	-3.25
21	67,229	-0.0039	0.1161	-8.77
63	65,813	-0.0087	0.1945	-11.54
84	65,243	-0.0091	0.2180	-10.70
126	63,841	-0.0134	0.2682	-12.66
252	58,979	-0.0340	0.3918	-21.10
504	50,525	-0.0347	0.6669	-11.68
756	42,629	-0.0284	0.9752	-6.00
Panel B: Only Accounts With More Than 10 Trades				
5	15,356	-0.0004	0.0242	-2.15[a]
21	15,352	-0.0012	0.0493	-2.99
63	15,333	-0.0035	0.0816	-5.26
84	15,321	-0.0042	0.0937	-5.49
126	15,277	-0.0061	0.1192	-6.37
252	15,101	-0.0179	0.1882	-11.67
504	14,598	-0.0093	0.3928	-2.85
756	13,900	-0.0185	0.6006	-3.63

[a] Significant at the 5% level.

Table 4.4: Trade Classifications

This table shows the percentage of trades classified as substitution trades. 'Days' is the window of days within which transactions must take place to be classified as substitution trades, and 'Maximum Buy/Sell Difference' is the largest allowed difference between the value of trades sold and the value of trades purchased. For example, 5.1% of transactions are classified as substitution trades when we require all transactions in a substitution trade to be within a two-day period and also that the total value of shares purchased and shares sold be within 2% of each other.

	Maximum Buy/Sell Difference				
Days	1%	2%	5%	10%	25%
1	2.9	5.0	9.0	13.2	19.9
2	3.0	5.1	9.5	14.1	21.6
3	3.1	5.4	10.1	14.9	23.1
5	3.2	5.6	10.7	15.9	25.1
10	3.6	6.3	12.1	18.5	30.1
15	3.8	6.8	13.1	20.2	33.1

Table 4.5: Post-Transaction Returns in U.S. Data

This table presents the results of the regression

$$R_{i,t} = \alpha + \beta_1 \cdot \text{Buy} + \beta_2 \cdot \text{Substitution} + \beta_3 \cdot \text{Buy} \times \text{Substitution} + \epsilon_{i,t},$$

where R denotes returns, Buy is a dummy variable that takes a value of one for purchases and zero for sales, and Substitution is a dummy variable that takes a value of one when a trade is classified as part of a substitution trade, and zero otherwise. 56% of transactions are purchases. A trade is classified as a substitution trade if all transactions occur within one trading day of each other, and the total value of shares purchased is within 2% of the total value of shares sold. Using this classification, 4.7% of transactions are classified as being part of a substitution trade. Stock returns are calculated as the increase in a cumulative return index (including all payouts) beginning the day after a transaction, and over the horizon specified in each column. Mutual fund returns are calculated from monthly returns reported in CRSP beginning in the month after the transaction. Standard errors, which are robust to correlation within stocks and mutual funds, are in parentheses. The number of clusters used in the robust standard error calculation is reported below the number of observations. $\beta_1 + \beta_3$ is the difference in the post-transaction return for a Substitution Buy and a Substitution Sell. $\beta_2 + \beta_3$ is the difference in return for a Substitution Buy and a Regular Buy. p-values for these F-tests are in square brackets. Statistical significance at the 10%, 5% and 1% levels is denoted by *, **, and ***, respectively.

	Post-Transaction Return Horizon			
	1 month	3 months	6 months	1 year
Panel A: No price restriction				
Constant	0.0140***	0.0358***	0.0672***	0.1602***
	(0.0006)	(0.0019)	(0.0039)	(0.0091)
Buy	-0.0022***	-0.0046***	-0.0083***	-0.0200***
	(0.0005)	(0.0012)	(0.0021)	(0.0040)
Substitution	-0.0038***	-0.0073***	-0.0062***	-0.0136***
	(0.0007)	(0.0013)	(0.0021)	(0.0038)
Buy×Substitution	0.0039***	0.0069***	0.0114***	0.0230***
	(0.0012)	(0.0018)	(0.0027)	(0.0055)
$\beta_2 + \beta_3$	0.0002	-0.0004	0.0051**	0.0094**
	[0.8609]	[0.7850]	[0.0167]	[0.0182]
$\beta_1 + \beta_3$	0.0018	0.0023	0.0031	0.0030
	[0.1557]	[0.2572]	[0.3305]	[0.6440]
Number of observations	2,386,439	2,376,519	2,358,667	2,317,528
Number of stocks/funds	11,578	11,506	11,389	11,061
Mean of dependent variable	0.0127	0.0330	0.0624	0.1487

Table 4.5, continued

	Post-Transaction Return Horizon			
	1 month	3 months	6 months	1 year
Panel B: Price > 5				
Constant	0.0131***	0.0349***	0.0664***	0.1618***
	(0.0006)	(0.0020)	(0.0041)	(0.0096)
Buy	-0.0019***	-0.0039***	-0.0073***	-0.0215***
	(0.0005)	(0.0012)	(0.0022)	(0.0041)
Substitution	-0.0031***	-0.0062***	-0.0048**	-0.0111***
	(0.0007)	(0.0013)	(0.0021)	(0.0037)
Buy×Substitution	0.0035***	0.0064***	0.0091***	0.0203***
	(0.0011)	(0.0017)	(0.0026)	(0.0052)
$\beta_2 + \beta_3$	0.0004	0.0003	0.0043**	0.0092**
	[0.6758]	[0.8443]	[0.0390]	[0.0176]
$\beta_1 + \beta_3$	0.0016	0.0026	0.0018	-0.0012
	[0.1639]	[0.1779]	[0.5438]	[0.8359]
Number of observations	2,169,367	2,163,009	2,150,619	2,120,179
Number of stocks/funds	10,104	10,069	10,015	9,839
Mean of dependent variable	0.0119	0.0325	0.0623	0.1494
Panel C: Risk-adjusted returns				
Constant	0.0157***	0.0364***	0.0695***	0.1626***
	(0.0007)	(0.0021)	(0.0043)	(0.0095)
Buy	-0.0030***	-0.0059***	-0.0102***	-0.0245***
	(0.0006)	(0.0014)	(0.0024)	(0.0048)
Substitution	-0.0050***	-0.0091***	-0.0118**	-0.0208***
	(0.0011)	(0.0019)	(0.0030)	(0.0050)
Buy×Substitution	0.0024***	0.0073***	0.0086***	0.0256***
	(0.0015)	(0.0024)	(0.0037)	(0.0095)
$\beta_2 + \beta_3$	-0.0026*	-0.0017	-0.0032	0.0048
	[0.0647]	[0.4200]	[0.2963]	[0.5342]
$\beta_1 + \beta_3$	-0.0006	0.0014	-0.0016	0.0011
	[0.7147]	[0.6074]	[0.6905]	[0.9155]
Number of observations	2,169,367	2,163,009	2,150,619	2,120,179
Number of stocks/funds	10,219	10,104	9,857	8,592
Mean of dependent variable	0.0140	0.0330	0.0637	0.1488

Table 4.6: Post-Transaction Returns in Finland Data

This table presents the results of the regression

$$R_{i,t} = \alpha + \beta_1 \cdot \text{Buy} + \beta_2 \cdot \text{Substitution} + \beta_3 \cdot \text{Buy} \times \text{Substitution} + \epsilon_{i,t},$$

where R denotes returns, Buy is a dummy variable that takes a value of one for purchases and zero for sales, and Substitution is a dummy variable that takes a value of one when a trade is classified as part of a substitution trade, and zero otherwise. 50% of transactions are purchases. A trade is classified as a substitution trade if all transactions occur within one trading day of each other, and the total value of shares purchased is within 2% of the total value of shares sold. Using this classification, 3.2% of transactions are classified as being part of a substitution trade. Stock returns are calculated as the increase in a cumulative return index (including all payouts) beginning the day after a transaction, and over the horizon specified in each column. Standard errors, which are robust to correlation within stocks and mutual funds, are in parentheses. The number of clusters used in the robust standard error calculation is reported below the number of observations. $\beta_1 + \beta_3$ is the difference in the post-transaction return for a Substitution Buy and a Substitution Sell. $\beta_2 + \beta_3$ is the difference in return for a Substitution Buy and a Regular Buy. p-values for these F-tests are in square brackets. Statistical significance at the 10%, 5% and 1% levels is denoted by *, **, and ***, respectively.

	Post-Transaction Return Horizon			
	1 month	3 months	6 months	1 year
Constant	0.0345***	0.1338***	0.1313***	0.0938***
	(0.0067)	(0.0212)	(0.0147)	(0.0337)
Buy	-0.0193***	-0.0787***	-0.1114***	-0.1378***
	(0.0080)	(0.0225)	(0.0187)	(0.0280)
Substitution	-0.0203***	-0.0744***	-0.1113**	-0.1495***
	(0.0038)	(0.0143)	(0.0160)	(0.0148)
Buy×Substitution	0.0217***	0.0845***	0.1180***	0.1409***
	(0.0059)	(0.0179)	(0.0183)	(0.0204)
$\beta_2 + \beta_3$	0.0014	0.0101	0.0067	-0.0086
	[0.0647]	[0.4200]	[0.2963]	[0.5342]
$\beta_1 + \beta_3$	0.0024	0.0058	0.0066	0.0031
	[0.7147]	[0.6074]	[0.6905]	[0.9155]
Number of observations	6,242,331	6,242,331	6,070,921	5,633,849
Number of stocks/funds	149	149	149	149
Mean of dependent variable	0.0248	0.0928	0.0732	0.0214

Table 4.7: Returns and Aggregate Buy Ratios

This table presents the results of a regression of stock returns in a given period (days or weeks in Panels A and B, respectively) on either the Buy Ratio formed from aggregate substitution trading, or the Buy Ratio formed from aggregate non-substitution (regular) trading:

$$R_{i,t+h} = \alpha + \beta \cdot BR_{i,t} + \epsilon_{i,t}, \quad \text{for } h = -4, \ldots, 4,$$

where BR denotes the Buy Ratio, defined as the ratio of the number of purchases to total number of trades. Separate regressions are run for each time horizon and the Buy Ratio constructed either from regular trades, or from substitution trades, is used, so each cell displays the results of a separate regression. Statistical significance at the 10%, 5% and 1% levels is denoted by *, **, and ***, respectively.

	Return Period Relative to Buy Ratio								
	-4	-3	-2	-1	0	1	2	3	4
	Panel A: Days								
Regular Buy Ratio	-0.0001 (0.0001)	-0.0002 (0.0001)**	-0.0001 (0.0001)	-0.0001 (0.0001)	0.0001 (0.0001)	-0.0003 (0.0001)***	-0.0002 (0.0001)**	-0.0003 (0.0001)***	-0.0002 (0.0001)**
Number of Observations	771,153	769,743	769,077	770,579	770,684	770,069	769,491	770,163	771,516
Informed Buy Ratio	0.0002 (0.0002)	0.0002 (0.0002)	0.0003 (0.0002)	0.0000 (0.0002)	0.0005 (0.0002)***	0.0002 (0.0002)	0.0003 (0.0002)*	0.0003 (0.0002)	0.0003 (0.0002)*
Number of Observations	177,512	177,167	176,984	177,628	177,311	177,042	177,193	177,170	177,546
	Panel B: Weeks								
Regular Buy Ratio	-0.0045 (0.0003)***	-0.0049 (0.0003)***	-0.0061 (0.0003)***	-0.0077 (0.0004)***	-0.0078 (0.0004)***	0.0006 (0.0003)*	0.0001 (0.0003)	0.0003 (0.0003)	-0.0000 (0.0003)
Number of Observations	455,671	456,528	457,383	458,322	459,444	459,698	459,089	458,502	457,867
Informed Buy Ratio	0.0001 (0.0006)	-0.0010 (0.0006)*	-0.0011 (0.0007)	0.0024 (0.0008)***	0.0111 (0.0010)***	0.0004 (0.0005)	-0.0007 (0.0005)	0.0009 (0.0005)*	-0.0000 (0.0005)
Number of Observations	98,001	98,161	98,310	98,469	98,745	98,838	98,751	98,665	98,581

Table 4.8: Returns and Aggregate Buy Ratios

This table presents the results of a regression of weekly stock returns on either the Buy Ratio formed from aggregate substitution trading, or the Buy Ratio formed from aggregate non-substitution (regular) trading:

$$R_i(t+h) = \alpha + \beta_1 \cdot BR_R(t) + \beta_2 \cdot BR_I(t) + \beta_3 R_i(t+h-1) + \epsilon_{i,t}, \quad \text{for } h = -3, \ldots, 3,$$

where BR denotes the Buy Ratio, defined as the ratio of the number of purchases to total number of trades. Separate regressions are run for each time horizon and the Buy Ratio constructed either from regular trades, or from substitution trades, is used, so each cell displays the results of a separate regression. Statistical significance at the 10%, 5% and 1% levels is denoted by *, **, and ***, respectively.

	-3	-2	-1	0	1	2	3
Regular buy ratio	-0.0120***	-0.0136***	-0.0223***	-0.0457***	0.0018	0.0025***	0.0017*
	(0.0014)	(0.0017)	(0.0023)	(0.0035)	(0.0013)	(0.0008)	(0.0009)
Substitution buy ratio	0.0012**	0.0011**	0.0036***	0.0039***	0.0019***	0.0008	0.0012**
	(0.0005)	(0.0005)	(0.0006)	(0.0008)	(0.0005)	(0.0005)	(0.0006)
Constant	0.0068***	0.0083***	0.0116***	0.0240***	-0.0010	-0.0011***	-0.0012**
	(0.0007)	(0.0009)	(0.0012)	(0.0020)	(0.0007)	(0.0003)	(0.0005)
Lag return	0.0508***	0.0734***	0.0898***	0.0642***	0.0662***	0.0209	0.0126
	(0.0113)	(0.0201)	(0.0228)	(0.0202)	(0.0240)	(0.0284)	(0.0109)
Number of observations	36,278	36,376	36,743	38,185	38,032	36,663	36,357

Bibliography

Bibliography

Abreu, D., and M. K. Brunnermeier, 2003, "Bubbles and Crashes," *Econometrica*, 71, 173–204.

Alexander, G. J., G. Cici, and S. Gibson, 2007, "Does Motivation Matter When Assessing Trade Performance? An Analysis of Mutual Funds," *Review of Financial Studies*, 20, 125–150.

Arrow, K., 1962, "The Ecomomic Implications of Learning By Doing," *Review of Economic Studies*, 29, 155–173.

Avramov, D., T. Chordia, and A. Goyal, 2006, "Liquidity and Autocorrelations in Individual Stock Returns," *Journal of Finance*, 61, 2365–2394.

Barber, B. M., and T. Odean, 2000, "Trading is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors," *Journal of Finance*, 55, 773–806.

Barber, B. M., and T. Odean, 2001, "Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment," *Quarterly Journal of Economics*, 116, 261–292.

Barber, B. M., and T. Odean, 2006, "All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors," Forthcoming, *Review of Financial Studies*.

Barber, B. M., T. Odean, and M. Strahilevitz, 2004, "Once Burned, Twice Shy: Naïve Learning, Counterfactuals and the Repurchase of Stocks Previously Sold," Working paper, UC Berkeley.

Barber, B. M., T. Odean, and N. Zhu, 2006, "Do Noise Traders Move Markets?," Working paper, UC Davis.

Barberis, N., and R. Thaler, 2005, "A Survey of Behavioral Finance," in Richard H. Thaler (ed.), *Advances in Behavioral Finance*, vol. 2, chap. 1, pp. 1–75, Princeton University Press.

Benartzi, S., 2001, "Excessive Extrapolation and the Allocation of 401(k) Accounts to Company Stock," *Journal of Finance*, 56, 1747–1764.

Black, F., 1986, "Noise," *Journal of Finance*, 41, 529–543.

Bolton, P., and C. Harris, 1999, "Strategic Experimentation," *Econometrica*, 67, 349–374.

Cai, F., and L. Zheng, 2004, "Institutional Trading and Stock Returns," *Finance Research Letters*, 1, 178–189.

Campbell, J. Y., 2006, "Household Finance," *Journal of Finance*, 61, 1553–1604.

Campbell, J. Y., S. J. Grossman, and J. Wang, 1993, "Trading Volume and Serial Correlation in Stock Returns," *Quarterly Journal of Economics*, 108, 905–939.

Campbell, J. Y., T. Ramadorai, and A. Schwartz, 2008, "Caught on Tape: Institutional Trading, Stock Returns, and Earnings Announcements," Working paper, Harvard University.

Carhart, M., 1997, "On Persistence in Mutual Fund Performance," *Journal of Finance*, 52, 57–82.

Chancellor, E., 2000, *Devil Take the Hindmost: A History of Financial Speculation*, Plume, New York, N.Y.

Choi, J., D. Laibson, B. Madrian, and A. Metrick, 2007, "Reinforcement Learning and Investor Behavior," Working paper, Yale University.

Chordia, T., and A. Subrahmanyam, 2004, "Order Imbalance and Individual Stock Returns: Theory and Evidence," *Journal of Financial Economics*, 72, 485–518.

Coval, J. D., D. A. Hirshleifer, and T. Shumway, 2005, "Can Individual Investors Beat the Market?," Working paper, University of Michigan.

Coval, J. D., and T. Shumway, 2005, "Do Behavioral Biases Affect Prices?," *Journal of Finance*, 60, 1–34.

Cox, D. R., 1972, "Regression Models and Life Tables," *Journal of the Royal Statistical Society, Series B*, 34, 187–220.

Cox, D. R., and D. Oakes, 1984, *Analysis of Survival Data*, Chapman and Hall, New York.

Daniel, K., M. Grinblatt, S. Titman, and R. Wermers, 1997, "Measuring Mutual Fund Performance with Characteristic Based Benchmarks," *Journal of Finance*, 52, 1035–1058.

De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann, 1991, "The Survival of Noise Traders in Financial Markets," *Journal of Business*, 64, 1–19.

Easley, D., R. F. Engle, M. O'Hara, and L. Wu, 2002, "Time-Varying Arrival Rates of Informed and Uninformed Trades," Working paper, Cornell University.

Easley, D., N. M. Kiefer, M. O'Hara, and J. B. Paperman, 1996, "Liquidity, Information, and Infrequently Traded Stocks," *Journal of Finance*, 51, 1405–1436.

Easley, D., and M. O'Hara, 1987, "Price, Trade Size, and Information in Securities Markets," *Journal of Financial Economics*, 19, 69–90.

Fama, E., and K. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 33, 3–56.

Fama, E., and J. MacBeth, 1973, "Risk, Return, and Equilibrium: Empirical Tests," *Journal of Political Economy*, 81, 607–636.

Feng, L., and M. Seasholes, 2005, "Do Investor Sophistication and Trading Experience Eliminate Behavioral Biases in Financial Markets?," *Review of Finance*, 9, 305–351.

Frazzini, A., 2006, "The Disposition Effect and Underreaction to News," *Journal of Finance*, 61, 2017–2046.

Friedman, M., 1953, "The Case for Flexible Exchange Rates," in *Essays in Positive Economics*, pp. 157–203, University of Chicago Press.

Genesove, D., and C. Mayer, 2001, "Loss-Aversion and Seller Behavior: Evidence from the Housing Market," *Quarterly Journal of Economics*, 116, 1233–1260.

Glosten, L. R., and P. R. Milgrom, 1985, "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14, 71–100.

Goetzmann, W. N., and A. Kumar, 2008, "Equity Portfolio Diversification," Forthcoming, *Review of Finance*.

Gompers, P. A., and A. Metrick, 2001, "Institutional Investors and Equity Prices," *Quarterly Journal of Economics*, 116, 229–259.

Greene, W., 2005, *Econometric Analysis*, Prentice Hall, 5th edn.

Greenwood, R., and S. Nagel, 2007, "Inexperienced Investors and Bubbles," Working paper, Stanford University.

Griffin, J. M., J. H. Harris, and S. Topaloglu, 2003, "The Dynamics of Institutional and Individual Trading," *Journal of Finance*, 58, 2285–2320.

Grinblatt, M., and B. Han, 2005, "Prospect Theory, Mental Accounting, and Momentum," *Journal of Financial Economics*, 78, 311–339.

Grinblatt, M., and M. Keloharju, 2000, "The Investment Behavior and Performance of Various Investor Types: A Study of Finland's Unique Data Set," *Journal of Financial Economics*, 55, 43–67.

Grinblatt, M., and M. Keloharju, 2001a, "How Distance, Language, and Culture Influence Stockholdings and Trades," *Journal of Finance*, 56, 1053–1073.

Grinblatt, M., and M. Keloharju, 2001b, "What Makes Investors Trade?," *Journal of Finance*, 56, 589–616.

Grossman, S., R. Kihlstrom, and L. Mirman, 1977, "A Bayesian Approach to the Production of Information and Learning by Doing," *Review of Economic Studies*, 44, 533–547.

Hamilton, J. D., 1994, *Time Series Analysis*, Princeton University Press, Princeton, NJ.

Hasbrouck, J., 1991, "Measuring the Information Content of Stock Trades," *Journal of Finance*, 46, 179–207.

Heckman, J., 1976, "The Common Structure of Statistical Models of Truncation, Sample Selection, and Limited Dependent Variables and a Simple Estimator for Such Models," *Annals of Economic and Social Measurement*, 5, 475–492.

Hvidkjaer, S., 2006, "Small Trades and the Cross-Section of Stock Returns," Working paper, University of Maryland.

Ivković, Z., C. Sialm, and S. J. Weisbenner, 2006, "Portfolio Concentration and the Performance of Individual Investors," Forthcoming, *Journal of Financial and Quantitative Analysis*.

Ivković, Z., and S. J. Weisbenner, 2005, "Local Does as Local Is: Information Content of the Geography of Individual Investors' Common Stock Investments," *Journal of Finance*, 60, 267–306.

Kaniel, R., G. Saar, and S. Titman, 2008, "Individual Investor Trading and Stock Returns," *Journal of Finance*, 63.

Keynes, J. M., 1936, *The General Theory of Employment, Interest and Money*, Macmillan.

Kyle, A. S., 1985, "Continuous Auctions and Insider Trading," *Econometrica*, 53, 1315–1335.

Lee, C., and M. Ready, 1991, "Inferring Trade Direction from Intraday Data," *Journal of Finance*, 46, 733–746.

Linnainmaa, J., 2006, "Learning from Experience," Working paper, University of Chicago.

Linnainmaa, J., 2007, "The Limit Order Effect," Working paper, University of Chicago.

List, J. A., 2003, "Does Market Experience Eliminate Market Anomalies?," *Quarterly Journal of Economics*, 118, 41–71.

Llorente, G., R. Michaely, G. Saar, and J. Wang, 2002, "Dynamic Volume-Return Relation of Individual Stocks," *Review of Financial Studies*, 15, 1005–1047.

Mahani, R., and D. Bernhardt, 2007, "Financial Speculators' Underperformance: Learning, Self-Selection, and Endogenous Liquidity," *Journal of Finance*, 62, 1313–1340.

Newey, W. K., and K. D. West, 1987, "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

Nicolosi, G., L. Peng, and N. Zhu, 2004, "Do Individual Investors Learn from Their Trading Experience?," Working paper, Yale University.

Nofsinger, J. R., and R. W. Sias, 1999, "Herding and Feedback Trading by Institutional and Individual Investors," *Journal of Finance*, 54, 2263–2295.

Obizhaeva, A., 2007, "Information vs. Liquidity: Evidence From Portfolio Transition Trades," Working paper, MIT.

Odean, T., 1998, "Are Investors Reluctant to Realize Their Losses?," *Journal of Finance*, 53, 1775–1798.

Odean, T., 1999, "Do Investors Trade Too Much?," *American Economic Review*, 89, 1279–1298.

Pastor, L., and R. F. Stambaugh, 2003, "Liquidity Risk and Expected Stock Returns," *Journal of Political Economy*, 111, 642–685.

Roll, R., 1984, "A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market," *Journal of Finance*, 39, 1127–1139.

San, G., 2007, "Who Gains More By Trading—Institutions or Individuals?," Working paper, Hebrew University of Jerusalem.

Scholes, M., and J. Williams, 1977, "Estimating Betas from Nonsynchronous Data," *Journal of Financial Economics*, 5, 309–327.

Seru, A., T. Shumway, and N. Stoffman, 2007, "Learning By Trading," Working paper, University of Michigan.

Shapira, Z., and I. Venezia, 2001, "Patterns of Behavior of Professionally Managed and Independent Investors," *Journal of Banking and Finance*, 25, 1573–1587.

Shefrin, H., and M. Statman, 1985, "The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence," *Journal of Finance*, 40, 777–790.

Shiller, R. J., 2005, *Irrational Exuberance (Second Edition)*, Princeton University Press, Princeton, N.J.

Shleifer, A., and R. W. Vishny, 1997, "The Limits of Arbitrage," *Journal of Finance*, 52, 35–55.

Shumway, T., and G. Wu, 2006, "Does Disposition Drive Momentum?," Working paper, University of Michigan.

Sias, R. W., L. T. Starks, and S. Titman, 2006, "Changes in Institutional Ownership and Stock Returns: Assessment and Methodology," *Journal of Business*, 79, 2869–2910.

Weber, M., and C. Camerer, 1998, "The Disposition Effect in Securities Trading: An Experimental Analysis," *Journal of Economic Behavior and Organization*, 33, 167–184.

Wermers, R., 1999, "Mutual Fund Herding and the Impact on Stock Prices," *Journal of Finance*, 54, 581–622.

Wooldridge, J. M., 1995, "Selection Corrections for Panel Data Models Under Conditional Mean Independence Assumptions," *Journal of Econometrics*, 68, 115–132.

Zingales, L., and A. Dyck, 2002, "The Bubble and the Media," in Peter K. Cornelius, and Bruce Kogut (eds.), *Corporate Governance and Capital Flows in a Global Economy*, Chap. 4, pp. 83–104, Oxford University Press, New York.